A reinforcement learning approach to generate equivalent circuit models for Electrochemical Impedance Spectroscopy

Ali Jaberi [*1] **Yonatan Kurniawan** [*2] **Robert Black** [1] **Shayan Mousavi M.** [1] **Kabir Verma** [3] **Zoya Sadighi** [1] **Jason Hattrick-Simpers** [2]

[*]Equal contribution [1]*Clean Energy Innovation Research Center, National Research Council Canada, Mississauga, ON, Canada* [2]*Department of Material Science and Engineering, University of Toronto, Toronto, ON, Canada* [3]*Cheriton School of Computer Science, University of Waterloo, Waterloo, ON, Canada*. Correspondence to: Jason Hattrick-Simpers jason.hattrick.simpers@utoronto.ca.

1. Introduction

Electrochemical impedance spectroscopy (EIS) is a powerful characterization technique for probing electrochemical systems such as batteries, fuel cells, and corrosion processes. A central challenge in EIS analysis is the construction of an appropriate equivalent circuit model (ECM) that faithfully represents the underlying physicochemical mechanisms. Traditionally, ECM selection and parameterization rely heavily on expert intuition and manual trial-and-error, limiting scalability and posing a bottleneck for accelerated materials discovery and self-driving laboratory workflows. While recent machine learning approaches have sought to automate EIS analysis, most existing methods frame the problem as supervised classification or regression [1], implicitly assuming fixed circuit topologies and thus restricting generalizability across systems.

In this work, we introduce a generative framework for ECM discovery based on reinforcement learning (RL), which formulates ECM construction as a sequential decision-making problem using reward signals. In this approach, an RL agent iteratively modifies circuit topology by applying targeted structural operations—such as adding, removing, or reconfiguring circuit elements—guided by a reward signal based on the quality of fit between the actual EIS and the ECM-predicted values, as well as the complexity of the circuits. By framing ECM discovery as an adaptive search over a combinatorial space of circuit architectures, this approach enables autonomous exploration and data-driven control of model complexity, without reliance on predefined circuit families. As a result, the proposed framework offers a flexible and interpretable pathway for automated ECM discovery from EIS data, bridging machine learning, electrochemical modeling, and self-driving experimental workflows.

2. Methods

Reinforcement learning (RL) naturally frames sequential decision-making using reward signals, where an agent selects actions to maximize rewards. For electrochemical impedance spectroscopy (EIS), we formulate the equivalent circuit model (ECM) discovery as a sequential construction problem: starting from an initial circuit, the agent iteratively modifies topology to improve agreement with impedance data while controlling complexity. This generative approach captures the combinatorial nature of ECM design without relying on predefined circuit classes. Learning is driven by agent trial-and-error without explicit ECM ground-truth labels, enabling an adaptive balance between model expressiveness and parsimony.

The RL environment evaluates candidate ECMs against EIS data. At each step, the agent observes a state comprising the target impedance spectrum and the current circuit topology, represented as a linear chromosome of circuit elements (R, C, L, P) and connection operators ("–" for series, "/" for parallel) [2]. This symbolic representation is deterministically transformed into a tree structure and interpreted as an electrical circuit. Actions correspond to symbolic mutations applied to the chromosome, specifying both the mutation location and the replacement symbol. After each mutation, the environment performs parameter fitting to infer optimal element values. The resulting goodness-of-fit provides the primary contribution to the reward. An additional term in the reward function is added to favor simpler circuit topologies when multiple circuits achieve comparable agreement with the impedance data. This reward-based formulation enables learning without access to ground-truth circuit labels.

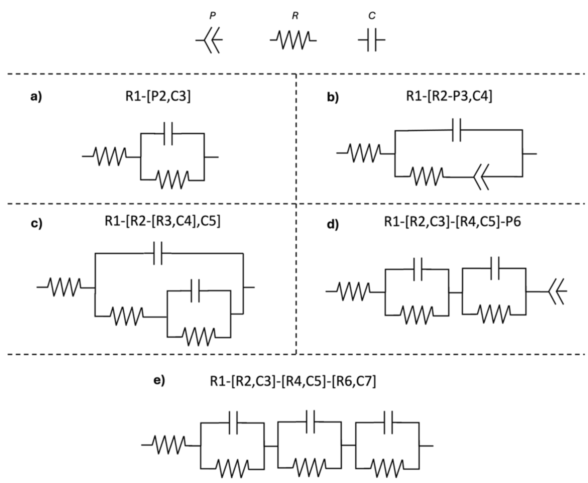

Fig. 1: Circuit topologies included in the training and test datasets.

To train and evaluate the agent, we generated a diverse dataset of EIS spectra adapted from literature [3], spanning multiple circuit topologies. The dataset includes single Randles circuits, Randles circuits augmented with Warburg elements, and multi-Randles configurations (two- and three-branch variants), as illustrated in Fig. 1. For each topology, 300 EIS spectra were generated by sampling circuit parameters from physically meaningful ranges, producing a wide variety of impedance responses. The data were split 90–10 for training and testing the RL agent, with the test set reserved for evaluating the agent's performance on unseen data.

3. Results

Table 1 summarizes the performance of the RL agent on the training set, a held-out test set, and an unseen experimental dataset. Performance is evaluated on a per-spectrum basis, where each case corresponds to an independent impedance dataset. We report the success rate, defined as the fraction of spectra for which the agent terminates with a valid equivalent circuit model, and the coefficient of determination (R^2) between the measured and fitted impedance responses. For R^2, averages are computed over all spectra as well as over successful cases only, since unsuccessful runs do not yield meaningful circuit fits.

Table 1: RL agent performance on the training set, test set, and experimental data.

	Train	Test	Exp.
Success rate	0.9963	1.0000	0.6897
Average R^2			
All data	0.9997	0.9999	0.6180
Successful	0.9999	0.9999	0.9977

On the training set, the agent achieves a success rate of 0.996 with near-perfect agreement with the data ($R^2 \approx 1.0$). Failures are confined to a small subset of the most complex circuit topology, namely circuit e shown in Fig. 1. Performance on the unseen test set closely matches that of the training data, yielding a perfect success rate and comparable R^2 values. Although these circuits are not encountered during training, they are drawn from the same design space, and the similar performance indicates consistent and reliable execution of the learned policy across distinct realizations.

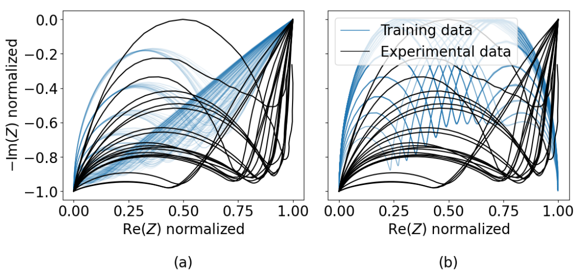

Fig. 2: Comparison of experimental EIS spectra with training set examples for two circuit topologies: (a) a single Randles circuit with a Warburg element and (b) a two-branch Randles configuration.

When applied to experimental EIS spectra, the agent is able to construct plausible equivalent circuit models, albeit with a reduced success rate. The experimental data exhibit characteristics associated with transitions between intermediate circuit topologies (Fig. 2), which are only indirectly represented in the training set. This result suggests that the agent has learned transferable structural features from the simulated data, enabling it to propose reasonable models even in the presence of topological ambiguity and experimental variability.

4. Conclusion

We present a generative RL framework for automated ECM discovery from EIS data, enabling flexible exploration of circuit topologies without predefined classes or ground-truth labels. The agent achieves near-perfect performance on both training and unseen test data drawn from the same design space, demonstrating stable and consistent execution of the learned policy. When applied to experimental spectra, the framework is able to propose physically plausible circuit models despite increased variability and topological ambiguity, highlighting its ability to extract transferable structural features from simulated data. Future improvements could incorporate active learning strategies, where uncertain or low-confidence experimental cases are iteratively selected for refinement through targeted simulation or expert feedback. Such an approach would allow the agent to adaptively expand its effective training distribution and further improve robustness when applied to complex, real-world EIS measurements.

Acknowledgments

This research was funded by the Government of Canada under the National Research Council Canada's Collaborative Science, Technology and Innovation Program (CSTIP)'s Critical Battery Materials Initiative (CBMI).

References

[1] Joachim Schaeffer, Paul Gasper, Esteban Garcia-Tamayo, Raymond Gasper, Masaki Adachi, Juan Pablo Gaviria-Cardona, Simon Montoya-Bedoya, Anoushka Bhutani, Andrew Schiek, Rhys Goodall, Rolf Findeisen, Richard D. Braatz, and Simon Engelke. Machine Learning Benchmarks for the Classification of Equivalent Circuit Models from Electrochemical Impedance Spectra. *Journal of The Electrochemical Society*, 170(6):060512, June 2023.

[2] Maxime Van Haeverbeke, Michiel Stock, and Bernard De Baets. Practical Equivalent Electrical Circuit Identification for Electrochemical Impedance Spectroscopy Analysis With Gene Expression Programming. *IEEE Transactions on Instrumentation and Measurement*, 70:1–12, 2021.

[3] Aleksei Makogon, Frederic Kanoufi, and Viacheslav Shkirskiy. Is Unsupervised Dimensionality

Reduction Sufficient to Decode the Complexities of Electrochemical Impedance Spectra? *ChemElectroChem*, 11(7):e202300738, 2024.