End-to-End Optimization for Battery Materials and Molecules by Combining Graph Neural Networks and Reinforcement Learning

Peter St. John,[1] Dave Biagioni,[1] Charles Tripp,[1] Jeffrey Law,[1] Erotokritos Skordilis,[1] Dmitry Duplyakin,[1] Struan Clark,[1] Robert Paton,[2] Shree Sowndarya S. V.,[2] Prashun Gorai,[3] Shubham Pandey,[3] Vladan Stevanovic,[3] Andrew Bray,[4] Troy Daley,[4] and John Meissner[4]

1 National Renewable Energy Laboratory
2 Colorado State University
3 Colorado School of Mines
4 Canvas Innovations

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated under Contract No. DE-AC36-08GO28308

Technical Report
NREL/TP-2700-90377
May 2025

End-to-End Optimization for Battery Materials and Molecules by Combining Graph Neural Networks and Reinforcement Learning

Peter St. John,[1] Dave Biagioni,[1] Charles Tripp,[1] Jeffrey Law,[1] Erotokritos Skordilis,[1] Dmitry Duplyakin,[1] Struan Clark,[1] Robert Paton,[2] Shree Sowndarya S. V.,[2] Prashun Gorai,[3] Shubham Pandey,[3] Vladan Stevanovic,[3] Andrew Bray,[4] Troy Daley,[4] and John Meissner[4]

1 National Renewable Energy Laboratory
2 Colorado State University
3 Colorado School of Mines
4 Canvas Innovations

Suggested Citation
St. John, Peter, Dave Biagioni, Charles Tripp, Jeffrey Law, Erotokritos Skordilis, Dmitry Duplyakin, Struan Clark, Robert Paton, Shree Sowndarya S. V., Prash. 2025. *End-to-End Optimization for Battery Materials and Molecules by Combining Graph Neural Networks and Reinforcement Learning*. Golden, CO: National Renewable Energy Laboratory. NREL/TP-2700-90377. https://www.nrel.gov/docs/fy25osti/90377.pdf.

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated under Contract No. DE-AC36-08GO28308

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Technical Report
NREL/TP-2700-90377
May 2025

National Renewable Energy Laboratory
15013 Denver West Parkway
Golden, CO 80401
303-275-3000 • www.nrel.gov

NOTICE

This work was authored in part by the National Renewable Energy Laboratory for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. The information, data, or work presented herein was funded by the Advanced Research Projects Agency-Energy (ARPA-E), U.S. Department of Energy, under Award Number DE-AR0001205. The views expressed herein do not necessarily represent the views of the DOE or the U.S. Government.

Cover photos (clockwise from left): Josh Bauer, NREL 61725; Visualization from the NREL Insight Center; Getty-181828180; Agata Bogucka, NREL 91683; Dennis Schroeder, NREL 51331; Werner Slocum, NREL 67842.

National Renewable Energy Lab
Final Scientific/Technical Report
End to End Optimization for Molecular and Atomistic Systems
WAS 19/CJ000/07/02

Award	DE-AR0001205
Sponsoring Agency	USDOE, Advanced Research Project Agency – Energy (ARPA-E)
Lead Recipient	National Renewable Energy Lab
Project Team Members	Peter St. John (NREL) Dave Biagioni (NREL) Charles Tripp (NREL) Jeffrey Law (NREL) Erotokritos Skordilis (NREL) Dmitry Duplyakin (NREL) Struan Clark (NREL) Robert Paton (CSU) Shree Sowndarya S. V. (CSU) Prashun Gorai (CSM) Shubham Pandey (CSM) Vladan Stevanovic (CSM) Andrew Bray (Canvas Innovations) Troy Daley (Canvas Innovations) John Meissner (Canvas Innovations)
Project Title	End-to-end optimization for battery materials and molecules by combining graph neural networks and reinforcement learning
Program Director	Dr. David Tew
Principal Investigator	Dr. Peter St. John
Contract Administrator	Kelly Harper
Date of Report	10/30/2022
Reporting Period	04/09/2020 – 09/30/2022

The information, data, or work presented herein was funded in part by the Advanced Research Projects Agency-Energy (ARPA-E), U.S. Department of Energy, under Award Number [DE-AR0001205]. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Please check the appropriate box:

☒ This Report contains no Protected Data.

☐ This Report contains Protected Data and the award allows data to be marked as protected. Refer to your Attachment 2 for guidance on how to appropriately mark Protected Data. The applicable notice is provided below:

☐ This Report contains SBIR/STTR Data and the award allows data to be marked as SBIR data. Refer to your Attachment 2 for guidance on how to appropriately mark SBIR Data. The applicable notice is provided below:

Table of Contents

Table of Figures/Tables

Public Executive Summary

The National Renewable Energy Laboratory (NREL), together with the Colorado School of Mines (CSM) and Colorado State University (CSU), has developed a machine learning-enhanced approach to the design of new battery materials. Currently, such materials are designed in part via numerous expensive high-fidelity computational simulations that predict the performance of a given composition. Even with computational screening tools, the vast landscape of possible molecular or crystal structures exceeds current and future computational capacity. Improving the efficiency by which new materials can be optimized will therefore disrupt the cost, risk, and time required to bring new energy solutions to the marketplace.

Predicting the properties of an organic molecule or periodic crystalline material given its structure has grown increasingly common. These approaches leverage large-scale computational and experimental databases and ML approaches such as graph neural networks. The inverse design problem of finding a material that possesses desired properties is substantially more challenging, since enumerating all valid material structures is not feasible. In this project, we leveraged recent success in reinforcement learning to efficiently navigate this high-dimensional search space. Just as algorithms can find the optimal chess moves from nearly limitless options, we train an approach to evolve a simple starting structure into a complex structure that possess the desired properties.

Our solution has been demonstrated by applying it to two related design application tasks for short- and long-term energy storage, respectively: (1) the design of solid-state ion conductors and (2) the design of organic redox-active materials. The project has resulted an open-source software library for material design, documented examples of applying the library to both organic and inorganic material optimization, and peer-reviewed publications detailing the data, computational models, and resulting candidate materials.

Acknowledgements

This work was authored by the National Renewable Energy Laboratory for the US Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. The information, data, or work presented herein was funded by the Advanced Research Projects Agency-Energy (ARPA-E), U.S. Department of Energy, under Award Number DE-AR0001205. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof. The US Government retains and the publisher, by accepting the article for publication, acknowledges that the US Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work or allow others to do so, for US Government purposes. RSP and SVSS acknowledge the RMACC Summit supercomputer, which is supported by the National Science Foundation (ACI-1532235 and ACI-1532236), the University of Colorado Boulder and Colorado State University, and the Extreme Science and Engineering Discovery Environment (XSEDE) through allocation TG-CHE180056.

Accomplishments and Objectives

This award allowed the National Renewable Energy Laboratory, Colorado School of Mines, and Colorado State University to demonstrate several key objectives. The focus of the project was on building software that leveraged artificial intelligence (AI) and machine learning (ML) techniques to dramatically speed up the search for high-performance energy materials.

Several tasks and milestones were laid out in Attachment 3, the Technical Milestones and Deliverables, at the beginning of the project. The actual performance against the stated milestones is summarized here:

Table 1. Key Milestones and Deliverables

Tasks	Milestones and Deliverables
Task 1: Graph Network Reinforcement Learning Algorithm Development	Q2: Developed a trainable reinforcement learning strategy for building organic molecules. Leveraged existing radical database to train a generative model that results in at least 50 novel molecules with high similarity to the training database. Novelty was assessed computationally using RDKit's SMILES canonicalization procedure, and similarity was validated using Tanimoto similarity of molecular fingerprints, ensuring all molecules have a similarity of at least 0.7 to existing calculations. **Actual Performance:** (10/08/20) A simple Monte Carlo Tree Search procedure was able to generate 286 unique SMILES strings that were unique from the database radicals yet had a Tanimoto similarity metric to at least one database radical of at least 0.7. This search could be completed in 20 minutes on a single compute node. Q3: Utilized the reinforcement learning approach developed in M1.1 to generate at least 30 new molecules meeting specific criteria for applicability in organic flow batteries. Molecules generated via reinforcement learning are evaluated using trained prediction algorithms from M4.2. New molecules have a predicted open shell voltage of at least 1V at 298K, and a predicted stability metric (from M4.1) superior to the established threshold, and an X–H bond dissociation energy <80 kcal/mol. (Joint with Tasks 2 & 4). Develop a working software prototype that documents the required computational steps and libraries involved. **Actual Performance:** (01/08/21) The developed reinforcement learning algorithm was used to find molecules meeting strict criteria for redox flow battery applications. A reward function incorporating the radical stability metric developed in this project, as well as constraints on the redox voltage for both ionization energy and electron affinity was developed and optimized. A total of 3,871 molecules were found that met the desired stability, redox, and BDE criteria specified in M1.2. These molecules were found with reasonable computational expense, requiring four hours of compute time on 10 CPU nodes and 1 GPU node on NREL's Eagle HPC system. The software prototype developed in this project has been released under an open-source BSD 3-Clause license and is available at https://github.com/nrel/rlmolecule/. Q5: Develop crystal structure RL algorithm. A reinforcement learning framework was developed for inorganic crystalline materials following

Tasks	Milestones and Deliverables
	the action space defined in M3.2. The method was demonstrated to generate at least 50 novel and stable crystal structures, confirmed via quantum mechanical density functional theory (DFT) calculations of the total energies (0K) of the structure of interest and its relevant competing phases. The formation enthalpies were calculated from DFT total energies using the fitted elemental-phase reference energies (FERE) methodology. Phase stability of a given structure was then assessed from the formation enthalpies of all the competing phases using the convex hull construction. DFT total energies were calculated using the plane-wave VASP software package. A high-throughput python framework (Pylada) was used for automating the calculation set up, job management, and analysis. **Actual Performance:** (10/08/21) The success of the RL framework was demonstrated by generating 382 structures predicted to have a negative decomposition energy with the surrogate model. Subsequent DFT validation confirmed 376 of these structures to be stable with respect to potential competing phases. Analysis of the generated structures indicated that while they were not present in known experimental databases, they appear chemically reasonable. Q6: Demonstrate success of RL on development of solid-state ion conductors. The RL framework (from M1.3) and predictive model from (M3.1) were used to generate at least 20 crystal structures that were predicted to be stable against decomposition into competing phases (including the anode / cathode chemistries chosen in M3.3) while containing the required elemental composition for ion (e. g. Li+) conduction. (Joint with Tasks 2 & 4). **Actual Performance:** (10/08/21) We validated the suitability of RL-generated candidates in a solid-state battery by computing their electrochemical stability windows. This property governs whether the material will decompose under the large voltage windows imposed on anode and cathode materials. At least 40 materials were shown to be stable at either metal anodes or high-voltage cathodes, or both.
Task 2: Software Engineering, Data Management, and Industry Engagement	Q1: Initial Commercialization Plan developed and communicated to ARPA-E that includes: (a) an explicit identification of the software and/or data that is expected to be a primary outcome of the project and subject to the U.S. Commercialization Plan and requirements set forth in Attachment 4, (b) targeted markets and customer segments; (c) customer value proposition, (d) competing alternatives, (e) potential strategic partnerships required for deploying the solutions (f) preliminary business model hypothesis, (g) IP strategy, (h) anticipated benefits to the US economy, and (i) a plan to report to ARPA-E the utilization of the software and data specified in the Commercialization Plan for a period of 10 years following its reporting. (This report is to include a description of any modifications made to the software or data and the manner in which such modifications are being commercialized.) Additional requirements that must be included in the Plan can be found in the clause in Attachment 2 entitled Intangible property (University awardee) or Rights in Data (for profit awardee). Any draft of the Commercialization Plan that has been approved by ARPA-E, supersedes any prior approved draft

Tasks	Milestones and Deliverables
	and is automatically incorporated by reference as part of this Agreement by the provisions of Attachment 2. **Actual Performance:** (07/08/20) An Initial Commercialization Plan was developed and communicated to ARPA-E that included the sections identified in the milestone text above. This plan outlines reporting requirements, targeted markets, and a preliminary business plan. This followed a commercial – open-source partnership, where NREL has lead the development of open-source software, while future industry partnerships would develop managed installations and deployments of the software within existing cheminformatic tools. Discussion with ARPA-E identified that additional work on discovering commercialization pathways should be included in future drafts, specifically about what specific materials modeling platforms the proposed software could be integrated with. The subsequently revised report also included additional discussion of competing alternatives and existing methodologies when compared to the proposed research. Q2: Software development plan developed and communicated to ARPA-E that includes: (a) a version control strategy (e.g. Github), (b) programming language & current version (e.g. Julia 1.1, Python 3.7), (c) major third-party/open source toolboxes (e.g. Scikit-learn, Tensorflow, JuMP), and (d) a software architecture and internal data flow plan. **Actual Performance:** (10/08/20) A software plan was developed and communicated to ARPA-E that included our plans for developing the RL framework and prediction methods via GitHub (and the locations for these packages when they are publicly released), the major programming language (Python) and ML framework (Tensorflow), as well as other packages that we will leverage for material and molecule optimization. Schematics for the software architecture and data flow plan were also produced, which shows how the internal components of the AlphaZero algorithm interact. Q3: Storage plan completed and communicated to ARPA-E that includes: (a) the anticipated size (e.g., number of terabytes), (b) location and (c) accessibility strategy (e.g. open source, limited access, proprietary). Data storage options should necessarily include the ability to easily update databases with newly generated computational results. **Actual Performance:** (01/08/21) A data storage plan was completed and communicated to ARPA-E that included the anticipated size, location, and accessibility strategy of data generated in this project. The main source of data preserved from this project are quantum mechanical calculations that were used to develop the criteria for optimized materials and confirm the top performing candidates.) Q6: Second Draft of U.S. Commercialization Plan and Updates to Data Acquisition/Storage and Software Development Plans. Commercialization plan updated to include: (a) preliminary cost/benefit analysis; (b) definition of potential launch markets, including critical performance and design features needed facilitate customer acceptance; (c) potential follow-on funding arrangements and (d) a detailed business model based on selected markets/competitive landscape. Additional requirements that must be included in the Plan

Tasks	Milestones and Deliverables
	can be found in the clause in Attachment 2 entitled Intangible property (University awardee) or Rights in Data (for profit awardee). Any draft of the Commercialization Plan that has been approved by ARPA-E, supersedes any prior approved draft and is automatically incorporated by reference as part of this Agreement by the provisions of Attachment 2. Data acquisition/storage and software development plans updated if necessary. **Actual Performance:** (10/08/21) An updated commercialization was created after several outreach interviews that provides further details on the cost/benefit proposition of the rlmolecule software, the main launch targets the software is intended to impact, possibilities for follow-on funding, and a detailed business model. The plan maintained and updated the required sections on reporting requirements and a full data & software list as specified by Attachment 2. The complete updated plan was attached to the Q5 quarterly report. Q8: Prepare documented release of RL/GNN package and deploy to cloud services. Release an open-source code repository containing the developed general-purpose libraries as well as well-documented code for the two application examples. Ensure the package is installable via common package managers (i.e., ananconda / pip) and document usage of the platform at least one cloud service provider such as AWS or GCP. **Actual Performance:** (04/08/22) A package for general-purpose graph optimization was developed and documented for a basic travelling salesman benchmark problem. This package is pip-installable, with documentation hosted on github pages. A manuscript detailing this open-source software package was submitted. A modification of the rlmolecule code to leverage this new library has also been completed. Documentation of running ray in AWS has been completed and is available at https://nrel.github.io/graph-env/cloud/tsp_on_aws.html. Q8: Final iteration of U.S. Commercialization and Formal Reporting of the Software Developed. Commercialization plan updated to include results of technical investigation, cost/benefit analysis, and market analysis to assess impact, value, and plans for further technological development. Specific updates include: (a) refined cost/benefit analysis; (b) results of market outreach, including formalization of partnerships with other entities to support next-stage development; (c) target performance characteristics of final prototype necessary to spur adoption; (d) revision/update of business model(s); (e) revision and/or expansion of IP plan; (f) propose organizational development requirements necessary to succeed in next-stage development activities. Report the completion of the software and data specified in the Commercialization Plan to ARPA-E per the procedures specified in Attachment 4 of the Award. Additional requirements that must be included in the Plan can be found in the clause in Attachment 2 entitled Intangible property (University awardee) or Rights in Data (for profit awardee). Any draft of the Commercialization Plan that has been approved by ARPA-E, supersedes any prior approved draft and is automatically incorporated by reference as part of this Agreement by the provisions of Attachment 2.

Tasks	Milestones and Deliverables
	Actual Performance: (09/30/22) A final reporting of the software and data created in this project was reported at the end of the project.
Task 3: Design Application to Electrolyte Interfaces for Solid-state Batteries	Q2: Developed a model to predict the formation enthalpy of a crystal structure with a mean absolute error of 50 meV/atom. Based on the predicted formation enthalpy of the structure of interest and its relevant competing phases, assessed whether the crystal structure will be stable using the convex hull construction. Crystal structures that are on or within 25 meV/atom from the convex hull, were deemed stable. Standard python packages were utilized to perform convex hull construction - inputs were stoichiometries and formation enthalpies of all competing phases. **Actual Performance:** (10/08/20) Crystal graph convolutional neural networks (CGCNNs) were trained to predict DFT total energies with an MAE of 36 meV/atom. Formation enthalpy was obtained from total energy through a simple linear transformation by using a tabulated library of elemental chemical potentials (constants). Therefore, the trained CGCNN model predicted the formation enthalpy with the same MAE as total energy (36 meV/atom). Q4: Formalize action space for building and manipulating crystal structures. A set of actions to incrementally build crystal structures, and for making predictions on the future reward of in-progress crystals, was formalized in a report for subsequent implementation in M1.3. **Actual Performance:** (04/08/21) An action space consisting of six steps was detailed in the Q2 quarterly report, which balanced the number of possible decisions at each step and the total number of decisions that must be made to construct a final structure. The search tree also grouped the most important decisions for an overall structure's properties at the higher-level branch points. This action space was randomly sampled to build a DFT energy training database for subsequent RL optimization. Q8: Confirm RL predictions for crystal properties with follow-up DFT calculations. Predicted values for optimal candidate materials from M1.3 were checked using slower density functional theory calculations. Example anode and cathode states were chosen to demonstrate the method. It was ensured through convex hull calculations against NREL's MatDB that at least 50% of these suggested materials were stable with respect to chosen representative anode and cathode materials. **Actual Performance:** (04/08/22) 46 novel MCTS-predicted crystal structures reported in M1.3 were checked with density functional theory for interfacial stability against electrodes. The electrochemical stability window was calculated using convex hull analysis, and more than 50% of the structures were found to be stable against high-voltage (4-5 V) cathodes. Several Li- and Zn-based, and few Na-based structures were found to be stable against the corresponding metal anodes.
Task 4: Design Application to Organic	Q1: Standard reduction potentials for at least 10,000 radicals in the BDE database were calculated using density functional theory at the M06-2X level of theory using Gaussian16 software. A metric to describe radical

Tasks	Milestones and Deliverables
Active Materials for Redox Flow Batteries	stability was derived from previous spin density delocalization calculations (included in the BDE database), and the ability of this metric to explain the stability of at least 6 well-known experimentally-verified stable radicals was demonstrated. A stability threshold and required prediction accuracy for the radical stability metric was also determined based on this experimental comparison. Radical stability metric and required prediction accuracy approved by the Program Director. **Actual Performance:** (07/08/20) 10,054 ionization energies and 13,487 electron affinities were calculated at the M06-2X / def2tzvp using 14,212 radicals taken from NREL's bond dissociation energy database. These totals only include calculations that were verified to have converged successfully and included checks on imaginary frequencies and molecular geometry. For instance, species which decomposed into different products upon the addition or removal of an electron were not included. These calculations were conducted using a water implicit solvent model. A radical stability metric was constructed from the fractional spin density of the unpaired electron and buried volume calculations. The maximum spin density fraction on any atom in the molecule was used as a molecular-level descriptor of radical delocalization, where smaller values indicate that the radical electron is spread over multiple atoms, while larger values (close to one) indicate the radical electron is localized on a single atom. Buried volume calculations were conducted on the atom in a molecule with the maximum spin density. These calculations describe the steric protection of radical through side groups that reduce the likelihood for radical reactivity, with higher buried volume indicating a more protected radical center. A stability metric was determined from these parameters by comparing seven known-stable radicals against those considered in the standard reduction potential database. As most organic radicals are short-lived, a classifier was fit to separate the likely short-lived radicals from the database from the 7 known to be stable experimentally. A linear support vector machine was able to fit a decision boundary that is surprisingly effective at separating known stable radicals from the computationally considered set. Q2: A graph neural network model was fit to directly predict the radical stability metric and standard reduction potential calculations generated in M4.1. Standard reduction potentials were predicted with a mean absolute error of at least 25 mV, and radical stability was predicted with an accuracy better than the baseline established in M4.1. Results were verified using a withheld test dataset of at least 1,000 radicals, calculated in M4.1. **Actual Performance:** (10/08/20) Separate graph neural network models were trained to predict radical stability and standard reduction potentials. In M4.1, we established a baseline accuracy in original units, of ±0.055 as acceptable for predictions of atomic spin density (spin density ranges from 0-1 in fractional units), while errors of ±2.75% as acceptable for buried volume predictions. The GNN trained to predict radical stability met these criteria, with an MAE of 0.009 for predictions of atomic spin density, and 0.73% for buried volume. A newly identified risk of this milestone was the accuracy limit of the standard reduction potentials identified above, which falls far below the standard 1 kcal/mol (43 mV) 'chemical accuracy' typically used as a gold standard for quantum

Tasks	Milestones and Deliverables
	chemistry or machine learning methods. In addition, the amount of training data required to reach chemical accuracy was underestimated in M4.1. To resolve this, nearly 10*x* more DFT calculations were performed and more stringent quality and convergence checks of DFT calculations were developed. A GNN model trained on redox potentials achieved a mean absolute error of 46.5 mV (1.07 kcal/mol) for ionization energies and 38.1 mV (0.88 kcal/mol) for electron affinities. While not meeting the 25mV threshold identified at the start of the project, these accuracies were more than sufficient for the optimization of stable radicals with desired redox properties in subsequent milestones. Q5: Verify RL-generated molecules with follow-up DFT calculations. Evaluation of the optimal candidates generated in M1.2 was performed with density functional theory (DFT). At least 50% of the molecules generated were confirmed (via DFT) to pass the standard reduction potential and radical stability thresholds established in M4.1. **Actual Performance:** (07/08/21) A selection of 208 top-scoring radicals generated via the RL algorithm was selected for follow-up confirmation via DFT. Of these, 138 radicals had successfully converged calculations for the radical, anion, and cation electronic forms (9 radical, 14 reduced, and 54 oxidized calculations terminated unsuccessfully or reached minima with bond dissociations or formations). Of the 199 radicals with converged DFT calculations, 126 (64.3%) had a stability threshold score greater than 85, well within the range of known-stable radicals. Of the 138 radicals with converged ionization energies and electron affinities, 115 (83.3%) had oxidation and reduction potentials within the water stability window and with a total cell potential greater than 1V. Combined, 74 radicals passed both DFT criteria, representing 53.6% of the converged calculations, although a smaller fraction (35.5%) of the entire generated set. Q7: Re-run RL algorithm to demonstrate improved candidates with subsequent iterations. The reinforcement learning algorithm (M1.1) and prediction models (M4.2) were retrained after including new calculations from M4.3. Newly generated compound predictions achieved the same standard reduction potential and radical stability targets (established in M4.1) as those in M1.2, but with a 20% lower molecular weight. **Actual Performance:** (01/08/22) Subsequent runs of the RL algorithm imposed a size limit of 12 heavy atoms, down from the original 15 heavy atoms used at the beginning of the project to demonstrate high-reward molecules. The surrogate objective functions have also been retrained using some of these previously generated high-reward molecules. The resulting molecules achieved similar performance targets with a 20% lower molecular weight.
Task 5: Project Management	Q1: Go-No/Go Milestone: Refine tasks and milestones for the work plan (if applicable) **Actual Performance:** (07/08/22) The milestone lists and task breakdown were unchanged.

Project Activities

This project focused on building software to enable inverse design of materials and molecules. The project used techniques from reinforcement learning, particularly the popular software RLLib, to enable high-throughput search over feasible structures. We hypothesized that by designing software that would allow the optimization of both inorganic and organic molecules, we would ensure that the developed software was general-purpose and suitable to industrial applications. The project resulted in several open-source software libraries for performing this optimization, papers describing the results published in top-tier research journals, and demonstrations of running the analysis on cloud-based systems.

This project made heavy use of computer modeling to provide training data for machine learning approaches and in building the machine learning models themselves. First principles quantum chemistry calculations were benchmarked against experimental measurements to determine the overall accuracy of the methods. Machine learning models were validated against held-out training data coming from quantum chemical simulations. All models considered in this project have been published in peer-reviewed publications or are currently undergoing peer review. A list of publications resulting from this project is presented in the following section.

This project received a no-cost extension of approximately 6 months, due in part to difficulty staffing the project during the COVID19 pandemic. The final milestones of the project (i.e., completion of the final software product) were delayed until the end of the FY22 fiscal year.

Project Outputs

A. Journal Articles

1. Sowndarya S. V., S., St. John, P. C., & Paton, R. S. (2021). A quantitative metric for organic radical stability and persistence using thermodynamic and kinetic features. *Chemical Science*, 12(39), 13158–13166. https://doi.org/10.1039/d1sc02770k
2. Pandey, S., Qu, J., Stevanović, V., St. John, P., & Gorai, P. (2021). Predicting energy and stability of known and hypothetical crystals using graph neural network. Patterns, 2(11), 100361. https://doi.org/10.1016/j.patter.2021.100361
3. S. V., S. S., Law, J. N., Tripp, C. E., Duplyakin, D., Skordilis, E., Biagioni, D., Paton, R. S., & St. John, P. C. (2022). Multi-objective goal-directed optimization of de novo stable organic radicals for aqueous redox flow batteries. Nature Machine Intelligence, 4(8), 720–730. https://doi.org/10.1038/s42256-022-00506-3
4. Biagioni, D., Tripp, C. E., Clark, S., Duplyakin, D., Law, J., & John, P. C. St. (2022). graphenv: a Python library for reinforcement learning on graph search spaces. Journal of Open Source Software, 7(77), 4621. https://doi.org/10.21105/joss.04621
5. Law, J. N., Pandey, S., Gorai, P., & St. John, P. C. (2022). Upper-Bound Energy Minimization to Search for Stable Functional Materials with Graph Neural Networks. https://doi.org/10.26434/chemrxiv-2022-sfxzl

B. Presentations

1. Rob Paton, 24th Annual ACS Green Chemistry & Engineering Conference (online) Jun 2020

2. Rob Paton, online seminar for GSK US/UK, Aug 2020

3. Rob Paton, online seminar for Genentech US, Nov 2020

4. Peter St. John, NREL SciML/AI Seminar series, Jan 2021

5. Peter St. John, DIFFERENTIATE Virtual Workshop on Reinforcement Learning, Jan 2021

6. Shubham Pandey, Jiaxing Qu, Vladan Stevanovic, Peter St. John, Prashun Gorai, Predicting Stability of Known and Hypothetical Crystal Structures using Graph Neural Network, Materials Research Society Spring Meeting 2021, Virtual, 2021

7. Peter St. John, Dave Biagioni, 2022 ARPA-E Energy Innovation Summit, May 2022.

8. Jeffrey Law, Designing Stable Battery Interface Materials Using Reinforcement Learning, Computational Materials Science and Engineering GRC, Aug 2022

9. Peter St. John, Goal-Directed Optimization of Stable Organic Radicals for Aqueous Redox Flow Batteries, Computational Materials Science and Engineering GRC, Aug 2022

10. Peter St. John, Reinforcement learning for battery material design, Telluride Science Research Workshop, Oct 2022

11. P. Gorai, Beyond the Known: Computational and Data-Enabled Discovery of Functional Materials, Materials for Humanity 2022 (Singapore MRS), Sep 2022.

C. Media Reports

Perspectives and Previews

1. Bartel, C. J. (2021). Data-centric approach to improve machine learning models for inorganic materials. Patterns, 2(11), 100382. https://doi.org/10.1016/j.patter.2021.100382

2. Cao, Y., Ser, C. T., Skreta, M., Jorner, K., Kusanda, N., & Aspuru-Guzik, A. (2022). Reinforcement learning supercharges redox flow batteries. Nature Machine Intelligence, 4(8), 667–668. https://doi.org/10.1038/s42256-022-00523-2

3. Hatcher, H. (2022). Game changer for batteries. Nature Reviews Chemistry, 6(9), 597–597. https://doi.org/10.1038/s41570-022-00425-2

News Articles

6. Dec. 2021 – "Machine Learning Method Could Speed the Search for New Battery Materials" https://www.nrel.gov/news/program/2021/machine-learning-method-could-speed-the-search-for-new-battery-materials.html
7. Sept. 2022 – "A molecular optimization framework to identify promising organic radicals for aqueous redox flow batteries" https://techxplore.com/news/2022-09-molecular-optimization-framework-radicals-aqueous.html

D. Software Records

8. Biagioni, D., Skordilis, E., Tripp, C., Duplyakin, D., St. John, P. "rlmolecule: A library for general-purpose material and molecular optimization using AlphaZero-style reinforcement learning." (2020) https://doi.org/10.11578/dc.20201221.3
9. Tripp, C., St. John, P., Biagioni, D., Clark, S., Duplyakin, D., Law, J. "Graph-Env." (2022) https://doi.org/10.11578/dc.20220622.3

E. Models and Databases

1. Law, J. (2022). jlaw9/upper-bound-energy-gnn (v0.1). Zenodo. https://doi.org/10.5281/zenodo.7089031
2. St. John, P., & Sowndarya S. V., S. (2022). pstjohn/redox-models (v0.0.2). Zenodo. https://doi.org/10.5281/zenodo.6390861
3. Sowndarya S. V., S.; St. John, P. (2021): Radical stability and redox potential calculations for 89,320 organic radicals in the water phase. figshare. Dataset. https://doi.org/10.6084/m9.figshare.14597556.v4

F. Status Reports

This project generated 9 quarterly reports, communicated directly to ARPA-E, and this final written report of the project's output.

Follow-On Funding

Additional funding committed or received from other sources (e.g., private investors, government agencies, nonprofits) after effective date of ARPA-E Award.

Table 2. Follow-On Funding Received

Source	Funds Committed or Received
DOE Basic Energy Sciences Phase II SBIR	$300,000.00