ARL-CR-0887 ● JAN 2026

Advancing Lithium-Mediated Electrochemical Ammonia Synthesis via High-Throughput Experimentation and Data Science

prepared by Michael Yusov, Ryan Jones, Gangsan Lee, and Karthish Manthiram
California Institute of Technology
1200 E. California Blvd., Pasadena, CA 91126

under contract W911NF-24-2-0154

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

ARL-CR-0887 ● JAN 2026

Advancing Lithium-Mediated Electrochemical Ammonia Synthesis via High-Throughput Experimentation and Data Science

prepared by Michael Yusov, Ryan Jones, Gangsan Lee, and Karthish Manthiram
California Institute of Technology
1200 E. California Blvd.
Pasadena, CA 91126

under contract W911NF-24-2-0154

REPORT DOCUMENTATION PAGE

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED	
		START DATE	END DATE
January 2026	Contractor Report	1 October 2024	30 September 2025

4. TITLE AND SUBTITLE
Advancing Lithium-Mediated Electrochemical Ammonia Synthesis via High-Throughput Experimentation and Data Science

5a. CONTRACT NUMBER	5b. GRANT NUMBER	5c. PROGRAM ELEMENT NUMBER
W911NF-24-2-0154		
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**

6. AUTHOR(S)
Michael Yusov, Ryan Jones, Gangsan Lee, and Karthish Manthiram

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
California Institute of Technology
1200 E. California Blvd.
Pasadena, CA 91126

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)	11. SPONSOR/MONITOR'S REPORT NUMBER(S)
DEVCOM Army Research Laboratory ATTN: RDRL-RLA-GC Adelphi, MD 20783-1138		ARL-CR-0887

12. DISTRIBUTION/AVAILABILITY STATEMENT
DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.

13. SUPPLEMENTARY NOTES
ORCID IDs: Michael Yusov: 0000-0002-0744-0469; Karthish Manthiram: 0000-0001-9260-3391

14. ABSTRACT
Electrochemical ammonia synthesis via the lithium-mediated nitrogen reduction reaction (LiNRR) offers a promising alternative to the energy- and carbon-intensive Haber–Bosch process, yet its development has been hindered by poor reproducibility, low selectivity, and limited mechanistic understanding. Here, we report on the establishment of a robust, high-throughput experimental framework that unlocks the use of data-driven optimization to accelerate LiNRR discovery. A specialized electrochemical cell was engineered to enable parallel experimentation, and careful analyses of the experimental procedure improved reproducibility. The resulting platform achieves approximately 2% relative standard deviation for repeat measurements, the highest reproducibility reported for LiNRR to date. Using these methods, we explored a multidimensional parameter space encompassing solvent composition, salt identity, proton donor content, and current density. Batch, multi-objective Bayesian optimization was shown to efficiently navigate this space, leading to the rediscovery of canonical high-performing electrolytes within just eight experiments and identification of novel formulations that improve both ammonia selectivity and production rate. These advances mark a critical step toward a quantitative, predictive approach to LiNRR and establish a foundation for future closed-loop studies of electrolyte design and optimization of electrocatalytic systems.

15. SUBJECT TERMS
Energy Sciences, lithium-mediated nitrogen reduction reaction, ammonia synthesis, Bayesian optimization, high-throughput platform, solid electrolyte interphase

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES
a. REPORT	b. ABSTRACT	C. THIS PAGE		
Unclassified	Unclassified	Unclassified	SAR	17

19a. NAME OF RESPONSIBLE PERSON	19b. PHONE NUMBER *(Include area code)*
Ana Alvarez	(520) 691-5235

STANDARD FORM 298 (REV. 5/2020)
Prescribed by ANSI Std. Z39.18

Contents

List of Figures

List of Tables

1. Introduction

Ammonia synthesis is critical for the production of fertilizer and many other nitrogen-containing substances. Moreover, ammonia is a promising fuel alternative due to its stability and energy density. However, ammonia is sourced primarily via the Haber–Bosch process, an environmentally hazardous operation that results in centralized production paired with global distribution networks. The carbon footprint and distribution of ammonia could be improved by exploring alternative pathways for production that utilize renewably sourced feedstocks under ambient conditions. Of such chemistries, electrochemical ammonia synthesis via the lithium-mediated nitrogen reduction reaction (LiNRR) has shown significant potential yet is underexplored. The efficiency and feasibility of this process are thus far limited by poor ammonia selectivity, low production rate, high cell potential, and poor electrolyte stability. Moreover, the discovery, understanding, and optimization of LiNRR are constrained by a lack of reliable, standardized data. These challenges can in part be attributed to deficiencies and differences in the current state-of-the-art cell designs used for analyzing LiNRR. In addition, LiNRR is hypersensitive to contamination; even air and water can cause significant perturbation in performance. Such sensitivities are believed to manifest themselves in the solid electrolyte interphase (SEI), a nanometer-scale passivation layer composed of electrolyte decomposition products. The SEI is believed to play a critical role in LiNRR systems due to its influence on nitrogen transport and proton availability. Therefore, when optimally designed, this tunable layer has the potential to significantly improve LiNRR selectivity and stability.

Despite its potential, electrolyte engineering for SEI design is difficult, not only due to the aforementioned sensitivity of this layer, but also due to the complexity of SEI formation. Modeling the effect of electrolyte composition on SEI composition and morphology remains a challenging task. In the absence of such models, intuition-driven electrolyte design is difficult. In this work, we present an alternative predictive approach for discovery and optimization of LiNRR. We first introduce a high throughput platform and experimental procedure that records record-breaking lows in the relative standard deviation of LiNRR performance metrics before pairing this platform with a multi-objective Bayesian optimization (BO) algorithm. We demonstrate preliminary trials of efficient optimization of LiNRR guided by BO and compare its performance to random exploration. This data-driven approach proves effective at predicting promising electrolyte compositions in a high-dimensional search space, demonstrating the effectiveness of using active learning for optimizing sensitive electrochemical systems.

2. Methods

2.1 Cell Design

Optimization over vast parameter spaces requires large amounts of data. To accelerate data acquisition, we designed a custom electrochemical cell capable of simultaneously performing independent experiments in parallel (Figure 1a). The cell is composed of four identical electrolysis chambers. Over several design iterations, we identified and eliminated key issues that caused electrolyte leaking and contamination between chambers. To accelerate prototyping and access geometries that cannot be machined using traditional methods, we implemented stereolithography-based 3D printing using Formlabs Rigid 10K resin, which was confirmed to exhibit chemical stability and low solvent swelling under LiNRR conditions. Moreover, comparison of 3D printed and machined polyether ether ketone/ polytetrafluoroethylene cell parts showed no measurable difference in ammonia selectivity, enabling rapid design cycles and validating printed materials for further experimentation.

The cell design incorporated gas-tight compression fittings and independent nitrogen and hydrogen flow paths. Stainless steel pressure plates fastened by stainless steel bolts ensured uniform sealing pressure. Cathodes and anodes were mounted in contact with copper current collectors and sealed with the help of O-rings, ensuring consistent electrode surface areas and facile replacement between runs (Figure 1b).

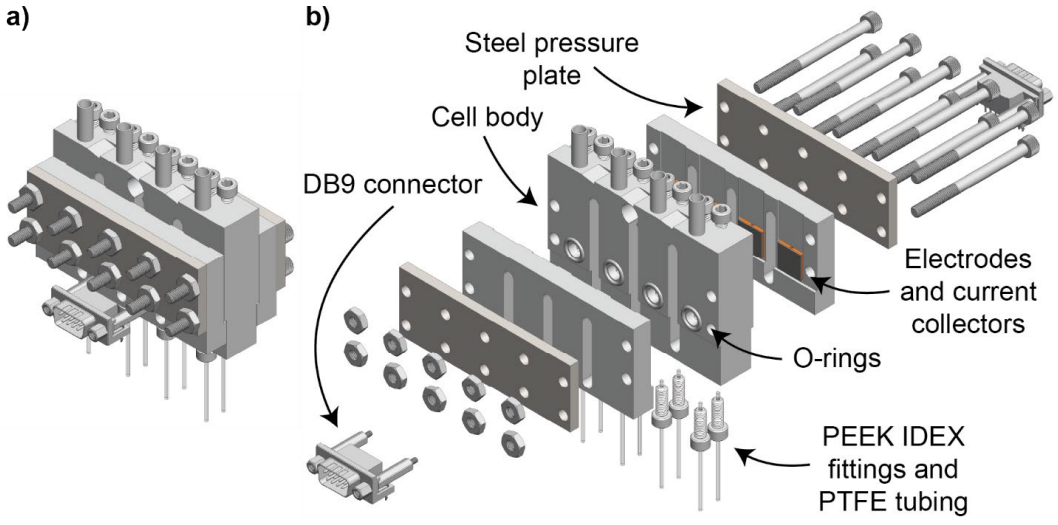

Figure 1. Detailed view of the electrochemical cell used for ammonia synthesis. a) Assembled 3D model highlighting the compact multi-chamber design for simultaneous experimentation. b) Exploded view illustrating individual components.

Gas lines equipped with mass-flow controllers (MFCs) and check valves were connected to inlet tubing. Nitrogen and hydrogen were passed through bubblers

containing dry diglyme (DG) or tetrahydrofuran (THF) immediately before entering the cell. This modular configuration enabled parallel experiments under strictly identical flow conditions, while maintaining isolated electrolyte chambers to avoid intermixing.

2.2 Experimental Protocol

Efficient and successful use of active learning requires high-quality data. For a system as sensitive as LiNRR, this calls for an experimental protocol that accounts for as many environmental or "hidden" variables as possible.

Inconsistencies in manual electrolyte preparation introduced subtle but systematic differences between experiments. To minimize contamination, particularly due to exposure to water and oxygen, all manipulations involving solvents, salts, and proton donors were conducted in an argon-filled glovebox. Solvents were dried over molecular sieves for at least 48 h, and salts were dried under vacuum at 120 °C for 48 h. Dried solvents were filtered through 0.45-µm polyvinyl difluoride syringe filters immediately prior to use to remove trace molecular sieve particulates. Prepared electrolytes were kept in the glovebox until ready for use. Any remaining errors due to inconsistencies in manual electrolyte preparation were captured by preparing independent batches of identical electrolytes, rather than single large electrolyte batches for use in multiple measurements.

Cell parts and electrodes were cleaned via sonication for 30 min sequentially in 10% H_2SO_4, acetone, isopropyl alcohol, and water. Cells were assembled after washing and dried overnight at 105 °C. Dry assembled cells were removed from the oven, immediately connected to N_2 and H_2 flow, and allowed to cool to room temperature. N_2 and H_2 flows were set to 18 and 2 sccm, respectively, and held constant throughout each experiment. After cooling, cell chambers were filled with electrolyte to a fixed volume of 1.5 mL per chamber. Cells were connected to a BioLogic potentiostat, and 30-min chronopotentiometry experiments began after 10 min of open circuit operation. Electrolytes were withdrawn immediately after electrolysis, and cell chambers were rinsed with water. Both the electrolyte and the water rinse solution were collected and combined for spectrophotometric quantification using the salicylate method.

Establishing control over gas delivery and electrolyte purity proved essential for reproducible operation. Chamber-to-chamber variations were kept minimal by regularly recalibrating MFCs with the help of a mass flow meter. Any potential bias due to small differences in inlet flows to different chambers was further minimized by randomizing the chamber order of conditions tested in any given batch of experiments.

2.3 Bayesian Optimization

The BO workflow comprised three phases: randomized initialization, quasi-random coverage via farthest-point sampling, and iterative multi-objective BO with batch candidate proposals. The design space consisted of seven continuous decision variables: current density (mA cm^{-2}), vol.% THF as solvent (the remainder being DG), $LiBF_4$ and LiTFSI concentrations (M), and 1-BuOH, 1-PrOH, and EtOH concentrations (M). Box bounds were enforced on each coordinate (current density 1–10 mA cm^{-2}; THF 0%–100%; salts 0–2.0 M; proton donors 0–0.15 M), together with linear inequality constraints on totals (salt sum 0.5–2.0 M; proton donor sum 0–0.15 M). A single global seed (12345) was used for reproducibility across NumPy, PyTorch, and Python's random module.

For randomized initialization, referred to as "Stage 1" herein, current density and vol.% THF were sampled from independent uniform distributions over their bounds. The total salt concentration was drawn uniformly and then partitioned into $LiBF_4$ and LiTFSI by a single uniform split fraction. The total proton-donor concentration was drawn uniformly and apportioned across 1-BuOH, 1-PrOH, and EtOH using a symmetric Dirichlet(1,1,1) distribution. This preserved the linear salt- and donor-sum constraints while exploring diverse salt and proton donor ratios. Each proposed composition was checked against box bounds prior to execution.

To improve coverage before BO, a quasi-random "Stage 2" used farthest-point sampling (FPS) in a reduced space defined by current density, %THF, total salt, and total proton donor. A candidate pool of 5,000 points was generated by independent uniform sampling over the reduced-space bounds. FPS proceeded greedily by selecting, at each step, the candidate that maximized the squared minimum Euclidean distance to the already selected set. To avoid over-selection near the edges—a common artifact of FPS—an edge-penalty factor was applied to the score: the distance term was multiplied by an "interiority" weight computed as a smooth function of the candidate's normalized distance to each bound. The selected reduced-space points were then lifted back to the full space by splitting salt with a uniform fraction and partitioning donors via Dirichlet(1,1,1), as in Stage 1, preserving the original linear constraints while achieving complementary coverage relative to the randomized batch.

For "Stage 3," multi-objective BO targeted maximization of ammonia Faradaic efficiency (FE) and production rate. Independent Gaussian-process (GP) surrogate models were fitted to each objective using BoTorch's SingleTaskGP with an input transform that normalized each of the seven inputs to [0,1] and an outcome transform that standardized each response. Exact GP hyperparameters were learned by maximizing the exact marginal log-likelihood (ExactMarginalLogLikelihood) with

GPyTorch, and a ModelListGP aggregated the two single-task models into a joint surrogate. Acquisition optimization used a q-Log Noisy Expected Hypervolume Improvement objective (qLogNEHVI), which directly maximizes the increase in dominated hypervolume in the presence of observation noise. A Sobol quasi-Monte-Carlo normal sampler was used to estimate the acquisition expectation, and the reference point for hypervolume was set to [0,0]. At each BO iteration, a batch of q = 4 candidates was generated by maximizing the acquisition under the original box bounds and the linear inequality constraints on salt- and donor-sums, using multi-start gradient-based optimization with SLSQP and BoTorch's batched initial-condition heuristics. Proposed candidates were executed at least in duplicate for each BO iteration, and newly acquired data were added to the training set for subsequent model fits.

3. Results

The experimental protocol was optimized based on observations of inconsistent performance when operating four chambers of a single cell in parallel under identical conditions. Notable changes included filtering solvents to remove molecular sieve particulates and ensuring proper MFC calibration. Some sources of variance remain unavoidable, such as from human error introduced during electrolyte preparation (Figure 2).

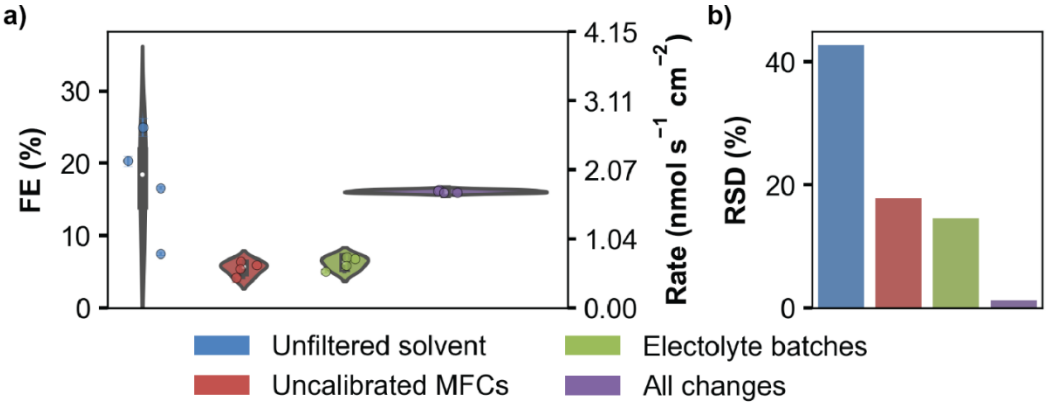

Figure 2. Reproducibility under different electrolyte preparation and experimental conditions. All experiments utilized 1 M $LiBF_4$, 0.1 M 1-BuOH in DG at –3 mA cm^{-2} with 18 sccm N_2 and 2 sccm H_2. "Unfiltered solvent" and "All changes" utilized 18 sccm N_2 and 2 sccm H_2, while "Uncalibrated MFCs" and "Electrolyte batches" utilized 10 sccm N_2 and 10 sccm H_2. a) Violin plots showing the distributions of four measurements, obtained under parallel operation. b) Relative standard deviations of the data plotted in (a).

The updated experimental procedure drastically improved reproducibility and reduced standard deviation in FE to below 1%, the highest reproducibility reported for LiNRR to the best of our knowledge. Under standard conditions of 1 M $LiBF_4$

and 0.1 M 1-BuOH in DG at –3 mA cm^{-2}, we obtain an average FE of 16.9 ± 0.7%, a rate of 1.75 ± 0.08 nmol s^{-1} cm^{-2}, and an energy efficiency of 2.27 ± 0.1% over 13 experiments. These results established a quantitative foundation for systematic data-driven electrolyte and parameter screening.

Screening and optimization proceeded in stages. Stages 1 and 2 illuminated the complex behavior of mixed electrolytes and provided a baseline training data set for subsequent BO (Figures 3a and 3b). Upon initializing the optimization algorithm, average performance began to increase, and the Pareto front expanded considerably, highlighting the effectiveness of BO (Figure 3c). At present, the Pareto front set contains just two experimental conditions that differ significantly, except for the choice and concentration of lithium salt (Figure 3d, Table 1).

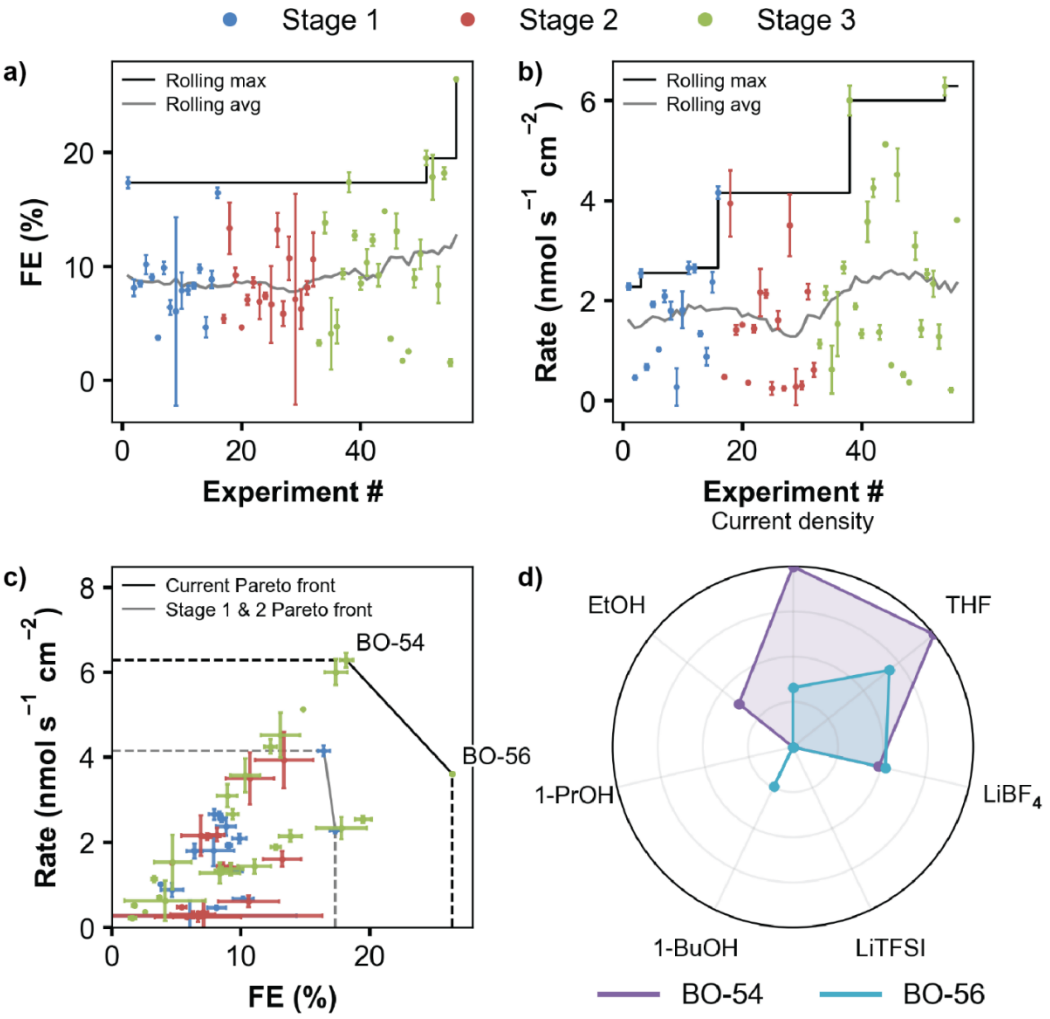

Figure 3. Results of screening and optimization. a,b) Progression of FE and rate with experiment number. Rolling averages were computed using a centered window with bin size 16. c) Summary of the obtained data, highlighting Pareto fronts at different stages. d) Star plot of conditions in the current Pareto front.

Table 1. Parameters of experimental conditions in the current Pareto front.

Pareto condition	Current density (mA cm^{-2})	THF (%)	$LiBF_4$ (M)	LiTFSI (M)	1-BuOH (M)	1-PrOH (M)	EtOH (M)
BO-54	10.00	100.00	0.975	0	0	0	0.0575
BO-56	3.956	68.239	1.05	0	0.0365	0	0

4. Discussion

Rolling averages for FE and rate did not change significantly with iteration until after initiating BO (Figures 3a and 3b). At this stage, the Pareto front expanded substantially in just six iterations (24 experiments), revealing conditions with both higher selectivity and rate than those observed during random screening (Figure 3c). In fact, BO discovered the most used electrolyte reported in LiNRR literature (~1 M $LiBF_4$, ~0.1 M EtOH in THF) in just two batches, or less than eight experiments. The algorithm also happened to pick the optimal current density for this electrolyte within the chosen constraints (–10 mA cm^{-2}) despite having zero knowledge of past findings apart from the data obtained during random screening. Successive iterations refined this discovery, identifying closely related compositions that offered marginally improved efficiency at reduced EtOH concentration (BO-54). The second Pareto condition (BO-56) differs from any reported in literature. The electrolyte contains both THF and DG and utilizes 1-BuOH instead of EtOH. This electrolyte yields higher FE but does so at a relatively low current density and thus has a lower rate than BO-54 (Figure 3d, Table 1). Further characterization, particularly of the SEI, is necessary to illuminate why and how these parameters improve selectivity.

Model performance can be further analyzed by considering the distribution of each parameter, along with that of the performance metrics (Figure 4). While Stages 1 and 2 show random and complementary distributions for THF, current density, salt concentration, and proton donor concentration, Stage 3 reveals preferences for solvent composition, current density, and salt concentration. In particular, the algorithm reveals a preference for high THF content, approximately 1 M $LiBF_4$, and current densities near the upper operational limit. The model remains uncertain about optimal proton donor concentrations, though the data shows a strong bias against the use of 1-PrOH.

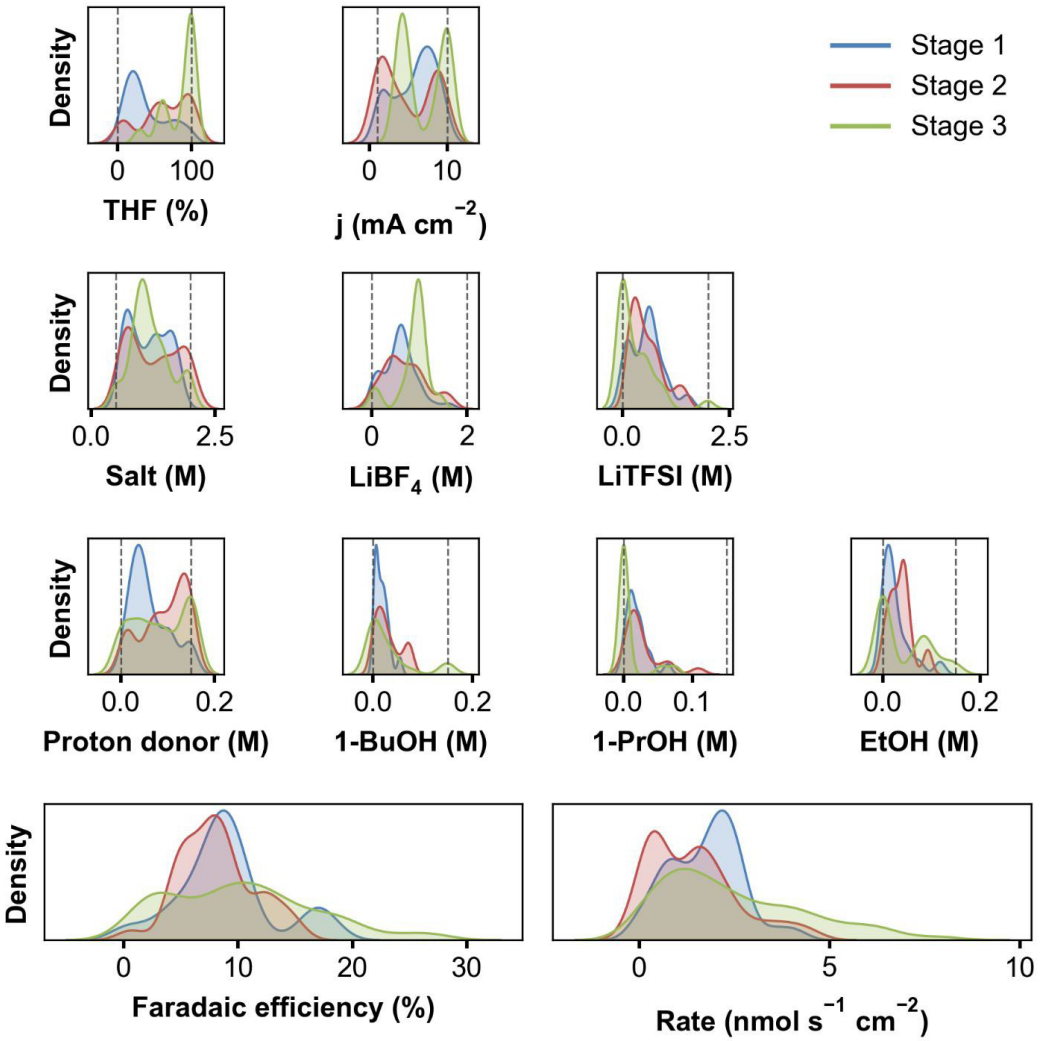

Figure 4. Kernel density distributions of parameters and target variables during the three stages of screening and optimization.

The bias towards the current density upper bound suggests that the original bounds could be limiting. We are planning an additional stage of optimization that will relax the upper bound on current density beyond 10 mA cm^{-2} by raising the lower bound of total salt concentration. This modification is designed to direct the acquisition toward compositions that retain high selectivity even at high current densities while maintaining the same linear-constraint framework and acquisition machinery.

The ongoing campaign will extend the dataset to approximately 100 experiments. Pareto-optimal conditions emerging from this process will then be tested in flow-cell architectures to compare performance and stability in a fully continuous reactor configuration with enhanced N_2 transport to current state-of-the-art conditions. Together, these efforts mark a transition from reproducibility, benchmarking, and screening to targeted discovery and potentially record-breaking enhancement of LiNRR performance.

5. Conclusion

This work has transformed lithium-mediated electrochemical ammonia synthesis from a fragile, often irreproducible process into one that is quantitatively robust, where even data-driven experimental frameworks can be applied. Through systematic reactor design, stringent control of experimental variables, and integration of data-driven optimization, we achieved record reproducibility and demonstrated the first practical implementation of active learning for LiNRR. Moreover, the progression from parallelized reactor engineering to statistically designed, model-guided experimentation helps establish a foundation for multi-objective closed-loop discovery in electrochemical systems beyond LiNRR.

The BO campaign has proved to be successful, underscoring the predictive power of data-driven approaches in electrocatalysis. Moving forward, we will extend this framework to additive discovery, using the same methodology to explore how molecular-level SEI structural changes govern macroscopic performance. Moreover, we will apply promising conditions in flow-cell geometries to bypass transport limitations, evaluate stability, and push performance beyond the current state-of-the-art. Further characterization using microscopic and spectroscopic characterization techniques will enhance our understanding of SEI engineering via electrolyte design. Collectively, these next steps will push LiNRR closer to practical viability while deepening our fundamental understanding of electrochemical nitrogen fixation.

List of Symbols, Abbreviations, and Acronyms

BO	Bayesian optimization
DG	diglyme
FE	Faradaic efficiency
FPS	farthest-point sampling
GP	Gaussian process
LiNRR	lithium-mediated nitrogen reduction reaction
MFC	mass-flow controller
SEI	solid electrolyte interphase
SLSQP	Sequential Least Squares Programming
THF	tetrahydrofuran

1 (PDF)	DEFENSE TECHNICAL INFORMATION CTR DTIC OCA
1 (PDF)	DEVCOM ARL FCDD RLB CI TECH LIB