

Integrated Machine Learning-Molecular Dynamics Framework for Electrolyte Property Prediction

N Rampal

January 2026

Energy and Environmental Science

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Integrated Machine Learning-Molecular Dynamics Framework for Electrolyte Property Prediction

Srikant Sagireddy^{*1,2,3}, Nikhil Rampal^{*1,2}, Stephen E. Weitzner^{1,2}, and Liwen F. Wan^{*1,2}

¹Materials Science Division, Lawrence Livermore National Laboratory, Livermore, California, 94550, United States

²Laboratory for Energy Applications for the Future (LEAF), Lawrence Livermore National Laboratory, Livermore, California, 94550, United States

³Department of Chemical Engineering, Stanford University, Stanford, California, 94305, United States

*srikant@stanford.edu, rampal1@llnl.gov, wan6@llnl.gov

Abstract

Electrochemical stability windows determine the operating range of battery electrolytes, yet accurate prediction remains challenging because stability emerges from statistical ensembles of local solvation environments rather than single ground-state molecular structures. Traditional density functional theory calculations on energy-minimized clusters cannot capture the thermal variations in local coordination environments and geometries that govern decomposition, while SMILES-based machine learning methods lack explicit representation of three-dimensional solvation structure and ion pairing. Here, we introduce a structure-aware machine learning framework that predicts frontier orbital energies (HOMO and LUMO) directly from molecular dynamics-sampled solvation configurations, achieving sub-0.6 eV accuracy at computational costs 3-4 orders of magnitude lower than first-principles methods. Across twelve representative battery electrolytes, we demonstrate that solvent-separated and contact ion pairs exhibit strong size- and local chemistry dependent electronic stability, with variations in coordination shifts of HOMO or LUMO level by 2-3 eV, and that extended solvation structure and partially desolvated environment further modulate stability by up to 3 eV. By encoding the statistical nature of electrochemical failure through ensemble sampling of explicit solvation geometries, our approach enables high-throughput screening and rational design of next-generation battery electrolytes with mechanistic understanding of structure-property relationships.

Introduction

Electrolyte design plays a central role in determining the performance, stability, and safety of next generation lithium-ion, sodium-ion, and multivalent batteries. Their ability to support efficient ion transport while maintaining electrochemical stability across wide operating voltages is essential for long-term device functionality. A key challenge in electrolyte design is the prediction of oxidative and reductive stability limits, which are often approximated using the energies of the highest occupied molecular orbital (HOMO), the lowest unoccupied molecular orbital (LUMO), and the corresponding HOMO-LUMO gap. While these quantities do not necessarily represent rigorous electrochemical stability limits, they provide computationally tractable and chemically interpretable surrogates that can be widely used for comparative screening of solvents, salts, and additives. As a result, frontier orbital descriptors have become a practical tool for guiding the selection of molecular components, such as salt-solvent combinations, and architectures for next-generation energy storage devices.

Experimental electrolyte screening is labor-intensive and economically challenging, and thus limits its practicality to explore broad chemical diversity. First-principles methods such as density functional theory (DFT) and *ab initio* molecular dynamics (AIMD) offer accurate predictions of electronic properties, but the computational expense becomes prohibitive for evaluating large molecular libraries or mapping the vast chemical design spaces [1]. These limitations have motivated the exploration of machine-learning (ML) approaches that can accelerate property prediction by several orders of magnitude compared to the conventional quantum mechanical methods.

Recent ML-driven studies have demonstrated the feasibility of screening electrolyte solvents and additives using learned models trained on DFT-derived electronic descriptors [2, 3]. Supervised and hybrid ML workflows have been used to identify candidate molecules with improved electrochemical stability and transport properties, often combining molecular descriptors with solvation-aware features [4, 1]. Notably, Zheng et al. [5] demonstrated that frontier-orbital descriptors derived from ML models compare to experimentally measured Coulombic efficiency and interfacial stability against lithium metal, enabling data-driven electrolyte screening guided by electronic structure metrics. However, current ML models face fundamental challenges in generalization, as these models are typically trained on datasets that represent limited families of electrolyte chemistries, leading to poor predictive performance on unseen solvents, salts, or coordination environments [6, 7]. Designing electrolytes that span carbonates, ethers, sulfones, phosphates, ionic liquids, fluorinated solvents, and reactive intermediates motivates the need for models that learn transferable chemical representations rather than correlations confined to narrow chemical domains.

Recent progress in ML suggests that generalizable predictive models arise from the integration of three key elements: dataset, ML modeling framework, and prediction task. The first element is a dataset that captures chemical diversity and structural realism. Electrolyte properties depend not only on molecular identity but also on solvation environments, coordination structure, ion pairing, and dynamic speciation [8, 9]. Datasets must therefore provide broad chemical coverage and include high-quality labels such as HOMO and LUMO energies, ionization potentials, solvation energies, and stability metrics that are computed with consistent and reproducible protocols. Foundation models provide a promising framework for addressing this challenge by learning transferable molecular and structural representations from large and chemically diverse datasets [10, 11, 12]. Their power arises from self-supervised representation learning, which enables the extraction of general chemical features such as orbital patterns, functional group behavior, intermolecular interaction motifs, and solvation structure without relying solely on expensive DFT labels [13, 14, 15]. By learning transferable molecular representations, these models provide a scalable route to transferability across wide regions of chemical space. In this context, foundation models act as representation learners rather than replacements for first-principles methods, enabling scalable screening while retaining a connection to underlying chemical physics.

The second element is a ML modeling framework capable of leveraging such datasets. This requires architectures that can represent both the compositional and structural complexity of electrolyte systems. These architectures may operate on molecular graphs, sequences, as in transformer-based models such as SMI-TED [16], or full three-dimensional atomic configurations, as in equivariant networks such as POS-EGNN [17], and include models ranging from message passing neural networks to transformer-based encoders and equivariant neural networks. The choice of representation governs the ability of a model to capture electronic structure, noncovalent interactions, and the influence of local environments. Scalable training

workflows that support multi-task and multi-fidelity learning are essential because electrolyte datasets often vary in size, accuracy, and computational cost. A robust architectural framework therefore provides the foundation for extracting transferable chemical information and for adapting models to diverse predictive tasks.

The third element is the prediction task itself. Properties that govern electrolyte performance include ionic conductivity, viscosity, solvation structure [18], and the electronic descriptors that define oxidative and reductive stability [13]. The HOMO, LUMO and gap have become practical surrogates for the electrochemical stability window (ESW) in many screening studies. Yet these properties do not originate solely from isolated molecular structure. They are strongly modulated by local chemical environments, including ion pairing, transient complexes, and fluctuating solvation environments. Such speciation can shift frontier orbital energies and modify redox limits in ways that gas-phase calculations do not capture. Accurate prediction therefore demands models that incorporate structural context and environment-dependent effects. These advances motivate a framework that moves beyond isolated-molecule descriptors toward environment-aware, transferable prediction of electrolyte electronic properties.

In this work, we introduce a structurally informed ML framework that predicts electrolyte frontier orbital energies from thermally sampled, solvated molecular environments, enabling transferable electronic property prediction across chemically diverse electrolyte systems. We filter chemically diverse datasets that capture realistic solvation environments, ion pairing behavior, and dynamic speciation, leveraging the foundational OMol25 database [11] to learn transferable molecular and structural representations. We incorporate physics-grounded structural information directly into the predictive framework by extracting thermalized molecular configurations from molecular dynamics (MD) simulations and using these structures as inputs to the equivariant graph-based neural network. By embedding local chemical context into the representation learning process, our approach bridges the gap between the efficiency of ML-based screening and the physical fidelity required for rational electrolyte design, thereby enabling rapid, transferable, and structurally informed screening of diverse electrolyte formulations for next-generation electrochemical devices.

Methods

Data acquisition

Our pre-training dataset originated from the open-source Open Molecules 2025 (OMol25) database [11]. This dataset comprises more than 100 million DFT calculations performed at the ω B97M-V/def2-TZVPD level of theory, representing a broad sampling of chemical space including small molecules, biomolecules, metal complexes, and electrolytes. Given our focus on battery electrolyte applications, we filtered the dataset to include only electrolyte structures, resulting in over 20 million electrolyte molecules. These 20 million pre-training structures provide a large and chemically diverse foundation for learning transferable representations of electronic structure in electrolyte-relevant molecules. We selected three key prediction targets for the pre-training phase: (1) HOMO energy level; (2) LUMO energy level; and (3) HOMO-LUMO gap.

To create a fine-tuning dataset tailored specifically to lithium-ion batteries, we sequentially filtered the pre-training dataset. We first identified structures containing lithium atoms, reducing the 20 million electrolyte structures to over 600,000 lithium-containing molecules, followed by removing all aqueous samples and under-coordinated lithium samples. We further retained only samples in which lithium coordination was between 3 and 10. These filtering steps produced a final fine-tuning dataset of approximately 24,000 structures representing chemically relevant Li-ion battery electrolytes.




Figure 1: ML training protocol. (A) Overview of workflow from data engineering to model training and verification. (B) Data processing. (C) Model training and inference.

Model architecture

Our model architecture is based on the IBM Foundation Models for Materials (FM4M) [19], particularly the POS-EGNN architecture [17]. The POS-EGNN model is built on GotenNet, a state-of-the-art geometric tensor network designed for efficient three-dimensional equivariant graph neural networks (EGNN) [20]. GotenNet addresses the fundamental trade-off between expressiveness and computational efficiency by employing geometric tensor representations that avoid computationally expensive alternatives, thereby maintaining practicality when scaling to large molecular pre-training datasets. Below we provide an overview of the model architecture.

Given a molecular system with N atoms, the model takes atomic numbers $\mathbf{z} \in \mathbb{R}^{N \times |\mathcal{Z}|}$ and three-dimensional coordinates $\mathbf{p} \in \mathbb{R}^{N \times 3}$ as inputs, where $|\mathcal{Z}|$ denotes the maximum atomic number, and then constructs a molecular graph by connecting atoms within a cutoff distance r_{cut} . Two types of node representations are maintained throughout its layers: invariant scalar features $\mathbf{h} \in \mathbb{R}^{d_{ne}}$ and equivariant steerable features $\tilde{\mathbf{X}}^{(l)} \in \mathbb{R}^{(1+2l) \times d_{ne}}$ for degree l , where d_{ne} represents the node embedding dimension. The edges are represented by invariant scalars $\mathbf{t}_{ij} \in \mathbb{R}^{d_{ed}}$ and tensors $\tilde{\mathbf{r}}_{ij} = \{\tilde{\mathbf{r}}^{(0)}, \tilde{\mathbf{r}}^{(1)}, \dots, \tilde{\mathbf{r}}^{(L_{\max})}\}$, where d_{ed} is the edge embedding dimension and each $\tilde{\mathbf{r}}^{(l)}$ encodes l -degree spherical harmonic projections of relative atomic positions, up to maximum degree $L_{\max}$. Information flows through the network via attention-driven message passing that explicitly incorporates spatial geometric information in each layer. The model outputs d_{ne} -dimensional feature vectors for each atom, capable of capturing both local atomic environments and global molecular structure.

For molecular-level property prediction, the node embeddings are aggregated through max pooling to

Table 1: Model architecture parameters for different model sizes.

Parameter	Small	Medium
Layers	4	6
$l_{\max}$	2	2
d_{ne}	256	384
d_{ed}	384	384
r_{cut} (Å)	6.0	6.0
Attention heads	8	8

obtain a holistic graph-level representation:

$$\mathbf{g} = \text{MaxPool}(\{\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_N\}). \quad (1)$$

This global molecular representation is then passed through a multi-layer perceptron prediction head to predict the HOMO/LUMO energies and HOMO-LUMO gaps. For atomic-level property prediction, like the site charges, the node embeddings $\mathbf{h}_i$ are directly passed through a separate prediction head.

Training strategy

We conducted pre-training of the POS-EGNN models using the AdamW optimizer with a constant learning rate of 1.0×10^{-4} and a batch size of 16, and trained 29 epochs over the complete 20 million electrolyte structures from the OMol25 dataset using PyTorch’s distributed data parallel (DDP) framework. The model parameters used for the different model sizes are provided in Table 1.

We then finetuned the model with identical architectural parameters but a reduced learning rate, 3.0×10^{-5} , for 10 epochs over the 24,000 selected Li-ion battery-relevant structures. The resulting fine-tuned model achieves enhanced performance on predicting relevant battery electrolyte solvation structures while retaining the broad generalization capabilities developed during pre-training.

At last, we implemented a multi-component loss function that balances prediction accuracy across all four targets. The total loss was defined as

$$\mathcal{L} = 10 \cdot \mathcal{L}_{\text{HOMO}} + 10 \cdot \mathcal{L}_{\text{LUMO}} + 1 \cdot \mathcal{L}_{\text{gap}} + 1 \cdot \mathcal{L}_{\text{SC}}, \quad (2)$$

where $\mathcal{L}_{\text{HOMO}}$, $\mathcal{L}_{\text{LUMO}}$, $\mathcal{L}_{\text{gap}}$, and $\mathcal{L}_{\text{SC}}$ represent the losses for HOMO energy, LUMO energy, HOMO-LUMO gap, and site charge, respectively. Each component loss was computed using the Huber loss function, which provides robustness to identify outliers. Since the HOMO-LUMO gap is mathematically derived as $\epsilon_{\text{gap}} = \epsilon_{\text{LUMO}} - \epsilon_{\text{HOMO}}$, accurate predictions of the individual orbital energies inherently ensure accurate gap predictions. We therefore assigned higher weights to the HOMO and LUMO losses to emphasize learning these fundamental electronic properties directly. The HOMO-LUMO gap loss, with a lower weight, was considered as a physical constraint that encourages the model to maintain meaningful relationship between HOMO and LUMO energies.

Results

Model performance

We first evaluate the predictive accuracy of our models on a held-out test set from the OMol25 battery electrolyte structures (See Section 2 of the Supporting Information (SI)). Fig. 2A shows the parity plots comparing model predictions against DFT-reference values for the HOMO, LUMO energies and the HOMO-LUMO gaps, annotated with the corresponding mean absolute error (MAE), root-mean-square error (RMSE), and Pearson correlation factor (ρ). The model achieves MAEs below 0.61 eV for all three targets, which is comparable to or smaller than the energy differences typically used to distinguish stable from unstable electrolyte solvents and additives. The high Pearson correlation coefficients ($\rho > 0.9$) observed for HOMO, LUMO, and HOMO-LUMO gap predictions imply that the model reliably preserves the relative ordering of frontier

orbital energies across chemically distinct species. As a result, it captures chemically meaningful trends in frontier energies across varied solvation environments, enabling robust ranking of candidate electrolytes. These trends are further elaborated in the following discussions.




Figure 2: Robustness analysis of ML model prediction. (A) Parity plots comparing the predicted HOMO, LUMO energies, and the HOMO-LUMO gap across the OMol25 Li-ion battery electrolyte test set for the small model. (B, C) t-SNE projection of Li-ion node embeddings color-coded by chemical composition (B) or solvation structure (C).

To assess whether the model learns chemically meaningful representations of local solvation environment beyond reproducing numerical trends, we analyze the internal representations learned by the model. Specifically, we extract node embeddings for lithium atoms from structures sampled during MD simulations (Section 1 of the SI) across a range of electrolyte systems to probe how local solvation environments are encoded. Fig. 2B and C show t-SNE projections of these Li-ion node embeddings. When color-coded by bulk electrolyte composition (Fig. 2B), the embeddings cluster according to solvent class, e.g. carbonate- or ether-based system, and anion identity, such as PF_6^- , FSI^- , or TFSI^- . When color-coded by local speciation (Fig. 2C), the embeddings separate into well-defined regions corresponding to freely solvated ions (blue), contact ion pairs (CIPs) (red), and larger aggregate structures (grey). This organization of the latent space indicates that the model encodes local coordination environments without explicit supervision on structural labels, capturing the fundamental and diverse chemical interactions that define local solvation environments.

Table 2: Small model predictions compared to DFT calculations across different electrolyte systems and local speciation.

Electrolyte	Speciation	HOMO [eV]			LUMO [eV]		
		Small	DFT	Error	Small	DFT	Error
LiFSI-G2	SSIP	-12.7	-12.7	0.0	-1.82	-1.70	0.12
	CIP	-9.70	-9.31	0.39	0.07	0.26	0.19
LiTFSI-G2	SSIP	-12.8	-12.8	0.0	-1.61	-1.64	-0.03
LiPF ₆ -EC	SSIP	-13.8	-13.2	0.6	-2.28	-1.56	0.72
	CIP	-11.1	-10.8	0.3	-0.50	0.11	0.61
LiPF ₆ -EMC	SSIP	-13.5	-13.3	0.2	-2.10	-1.75	0.35
	CIP	-11.8	-11.2	0.6	-0.72	0.19	0.91
LiPF ₆ -EC-EMC	SSIP	-10.8	-6.89	3.91	-4.10	-5.49	-1.39
	CIP	-9.54	-4.52	5.02	-3.45	-1.96	1.49
LiTFSI-EC-PC	SSIP	-13.1	-12.9	0.2	-1.55	-1.38	0.17
	CIP	-10.3	-9.79	0.51	0.01	0.20	0.19
LiFSI-DEE	SSIP	-13.8	-13.5	0.3	-2.20	-1.83	0.37
	CIP	-9.19	-8.97	0.22	-0.7	-0.09	0.61
LiTFSI-DEE	SSIP	-13.8	-13.5	0.3	-2.14	-1.82	0.32
	CIP	-10.8	-10.6	0.2	0.13	0.26	0.13
LiFSI-DPE	CIP	-9.99	-10.0	-0.01	0.55	0.44	-0.11
LiTFSI-DEE-DPE	SSIP	-13.2	-13.0	0.2	-1.96	-1.80	0.16
	CIP	-9.47	-8.86	0.61	0.03	0.10	0.07
LiTFSI-DOL-DME	SSIP	-7.36	-11.1	-3.74	-1.69	-4.08	-2.39
LiFSI-BFE	SSIP	-14.4	-14.2	0.2	-2.10	-1.87	0.23
	CIP	-9.05	-8.81	0.24	-0.15	-0.14	0.01
LiPF ₆ -BFE	SSIP	-14.1	-14.0	0.1	-1.94	-1.82	0.12
	CIP	-11.4	-11.0	0.4	0.23	0.22	-0.01
LiFSI-BTFE	CIP	-10.0	-10.1	-0.1	0.51	0.17	-0.34
LiPF ₆ -BTFE	SSIP	-13.9	-14.1	-0.2	-1.63	-1.88	-0.25
	CIP	-12.1	-12.3	-0.2	-0.12	-0.10	0.02

While the latent space analysis demonstrates that the model effectively captures local coordination physics, predictive accuracy must be quantitatively benchmarked against first-principles energies. To this end, we perform DFT single-point energy calculations at the ω B97M-V/def2-TZVPD level of theory on representative solvation structures extracted from the MD simulations (Section 3 of the SI). Table 2 compares the predictions from the small model against these DFT reference values across various electrolyte combinations, spanning carbonates, glymes, ethers, and fluorinated solvents with FSI⁻, TFSI⁻, and PF₆⁻ anions. For the majority of systems, the model achieves excellent agreement with DFT, with HOMO energy errors typically below 0.6 eV and LUMO errors below 0.4 eV. However, two systems exhibit notably larger deviations: LiPF₆-EC-EMC and LiTFSI-DOL-DME. These outliers share a common characteristic: they involve electrolyte mixtures in which multiple solvent species compete for interactions with lithium ions. We speculate that these errors arise from the model’s limited exposure to mixed-solvent coordination environments during pre-training, which may hinder its ability to fully capture complex solvent interactions in these specific systems. Still, promisingly, the model maintains high accuracy for most of the mixed or functionalized systems where the chemical components are structurally similar. For example, systems like LiTFSI-DEE-DPE (with errors below 0.61 eV) and LiTFSI-EC-PC (with errors below 0.51 eV) demonstrate excellent performance. Overall, these quantitative benchmark tests suggest that the model excels at interpolating within similar chemical spaces, but it may struggle to fully capture environments that combine

structurally distinct solvent classes, particularly those that were underrepresented in the pre-training data. Improving predictions for heterogeneous mixed-solvent systems would likely require augmenting the training dataset with more diverse multi-component coordination environments. Despite these current limitations, we believe the model remains robust and capable of providing quantitatively accurate frontier orbital predictions, which makes it suitable for comparative screening and rational electrolyte design. We leverage these capabilities in the following sections to analyze structure–property relationships, highlighting the model’s potential for guiding future electrolyte development.

Impact of local solvation environment on frontier orbital energies

Having established that the model both qualitatively and quantitatively predicts frontier orbital energies and encodes chemically meaningful representations of local solvation environments, we turn to a central question in electrolyte design: how local solvation structure influences electrochemical stability in practical battery electrolytes. To address this question, we begin by analyzing the predicted ESWs for a set of Li-ion battery electrolytes. The potentials of the cathodic and anodic limits, V_{CL} and V_{AL} , for single electron transfer can be readily extracted from the predicted HOMO and LUMO energies, via the relation

$$V_{\text{CL}} = -\frac{\epsilon_{\text{LUMO}}}{e} \quad (3)$$

and

$$V_{\text{AL}} = -\frac{\epsilon_{\text{HOMO}}}{e}, \quad (4)$$

where ϵ_{LUMO} and ϵ_{HOMO} are the energies of the predicted LUMO and HOMO levels, and e is the electron charge.

Based on the speciation, e.g., freely solvated ions, CIPs, SSIPs, etc. extracted from MD simulations of twelve different electrolyte systems, we predict their ESW using our benchmarked POS-EGNN model. The ensemble-resolved cathodic and anodic limits, defined as

$$V_{\text{CL}}^{\text{electrolyte}} = \max(V_{\text{CL,SSIP}}, V_{\text{CL,CIP}}) \quad (5)$$

and

$$V_{\text{AL}}^{\text{electrolyte}} = \min(V_{\text{AL,SSIP}}, V_{\text{AL,CIP}}), \quad (6)$$

are calculated and rendered in Fig. 3A. The colored bars indicate the regions in which the electrolyte is predicted to be stable, while the error bars denote the standard deviation of the electrochemical stability windows (ESWs) arising from dynamical variations in solvation structures sampled during MD simulations. As evident in Fig. 3A, substantial differences in ESWs across electrolyte chemistries, both in the absolute width of the window and in the relative positions of the cathodic and anodic limits. Even within the same solvent class, changes in anion identity or solvent mixing lead to systematic shifts in stability limits, underscoring strong sensitivity of frontier orbital energies to local solvation structures. In the following sections, we discuss the role of individual attributes of solvation environments to the predicted ESWs.




Figure 3: Correlation analyses. (A) Electrochemical stability windows of relevant Li-ion battery electrolytes. Colored bars denote the potential ranges over which the electrolyte is predicted to be stable, with the left and right edges corresponding to the cathodic and anodic limits, respectively. The vertical dashed black line marks the reference lithium metal Fermi energy at -2.91 eV, while the error bars reflect variations in the predicted HOMO and LUMO levels of solvated species arising from different coordination environments. (B) Histograms of HOMO, LUMO, and HOMO-LUMO gap for solvation structures color-coded by solvent class. (C) Representative snapshots of freely solvated Li-ion in DEE, extracted from MD simulation. (D) Representative snapshot of CIP configuration (Li-FSI in DEE). (E) Representative snapshot of secondary solvation shell of Li-ion in DEE. (F, G) Variations in HOMO, LUMO, and HOMO-LUMO gap as a function of the solvation shell size and composition, where **P** refers freely solvated Li-ion (primary solvation), **S** denotes secondary solvation environment and **S + A** represents secondary solvation with an anion species included.

Role of solvent-chemistry

As shown in Figure 3A, solvent chemistry has a significant impact on the predicted anodic and cathodic limits, as highlighted for LiFSI. In Diglyme (G2), the anodic limit is approximately 8.81 V, and the cathodic limit is around 1.75 V, resulting in a ESW of about 7.06 V. When G2 is replaced with ether solvents

like Diethyl ether (DEE), the anodic limit increases by around 0.87 V, leading to a slightly wider ESW. Further increasing the functional-chain length of the solvent molecule, such as with Dipropyl ether (DPE), extends the anodic limit to around 10.08 V, but with a significantly lowered cathodic limit (~ -0.45 V), effectively widening the overall ESW by more than 2 eV. When ether molecules are functionalized with fluorine, intermediate behaviors are observed. For example, LiFSI in bis(2-fluoroethyl) ether (BFE) exhibits a cathodic limit near 1.83 V and an anodic limit around 9.53 V. LiFSI in bis(2,2,2-trifluoroethyl) ether (BTFE) shows anodic and cathodic limits similar to those of DPE. Overall, for the same FSI anion, solvent choice can lead to variations of up to 2.0 V in both oxidation and reduction limits, demonstrating the significant influence of solvent chemistry on ESWs.

The strong solvent effects are further highlighted in the distribution analysis of HOMO, LUMO, and HOMO–LUMO gap values for all solvation structures extracted from MD sampling. As shown in Figure 3B, the distributions of HOMO and LUMO energies span a wide range, with clear separation based on solvent class. Glyme-based electrolytes exhibit the most stable HOMO and LUMO levels, with narrow distributions, suggesting they are less sensitive to geometrical variations in local solvation structures. In contrast, ether and fluorinated ether systems show broader distributions, indicating a stronger dependence of electrochemical stability on speciation and local solvation environments. Interestingly, carbonate mixtures exhibit the lowest-lying LUMO states, extending down to about -6 eV, which suggests increased susceptibility to reduction. The combined shifts in HOMO and LUMO levels are reflected in the HOMO–LUMO gap distributions, which generally cluster between ~ 8 and ~ 12 eV for most solvent systems. Overall, we observe the following trend for stability: carbonate mixtures < ethers, fluorinated ethers < carbonate < ether mixtures < glymes. The wide distribution of HOMO–LUMO gap values in fluorinated ethers suggests that these systems offer significant tunability through specific functionalization. This tunability presents a practical opportunity to design electrolyte systems with enhanced electrochemical stability, further expanding the range of viable electrolyte formulations for next-generation batteries.

Impact of anion interactions

While solvent chemistry establishes the foundational stability framework, the identity and coordination behavior of anion species can further revise the electrochemical stability limits of the electrolyte. The critical distinction lies not merely in anion identity, but also in how different anions modulate the population and electronic properties of distinct solvation environments, particularly the balance between CIPs and solvent-separated ion pairs (SSIPs).

Systematic comparison of LiPF₆, LiFSI and LiTFSI across different solvent systems reveals striking anion-dependent variations in both oxidation and reduction limits. In fluorinated ether systems, this pattern is particularly pronounced, where LiPF₆ in BFE exhibits an anodic limit of 11.04 ± 0.30 V compared to LiFSI in BFE at 9.53 ± 0.36 V, demonstrating a 1.5 V improvement in oxidative stability. Similarly, in BTFE solvent, LiPF₆ achieves 11.49 ± 0.39 V versus LiFSI’s 9.88 ± 0.33 V anodic limit. This consistent trend across different fluorinated ethers confirms LiPF₆[−] based systems provide superior oxidative resistance compared to sulfonyl-based salts. However, cathodic stability shows the opposite behavior. LiFSI systems consistently demonstrate enhanced reduction resistance compared to the LiPF₆ counterparts. In BFE solvent, both salts exhibit similar cathodic limits (1.83 V), but in BTFE, LiFSI achieves remarkable cathodic stability with a limit of -0.28 ± 0.24 V compared to LiPF₆’s 1.63 V. This 1.9 V difference represents one of the largest anion-induced cathodic shifts observed across all systems studied, highlighting the profound impact of anion identity on reduction processes.

The most dramatic anion effects appears in glyme-based electrolytes, where LiTFSI-G2 achieves an extraordinary anodic limit of 12.85 ± 0.13 V, nearly 4 V higher than LiFSI-G2’s 8.81 ± 0.86 V. This significant difference translates to stability windows of 11.11 V and 7.06 V, respectively, leading to a 57% increase in ESW. Remarkably, both systems exhibit nearly identical cathodic limits (1.74-1.75 V), meaning the extended ESW for LiTFSI-G2 is entirely due to enhanced oxidative resistance from the TFSI[−] anion. The structural flexibility of TFSI[−] introduces additional complexity through conformational effects. Unlike rigid anions, TFSI[−] can adopt multiple geometries around lithium, creating dynamic equilibria between different CIP structures. This conformational diversity is reflected in the broader error ranges observed for some TFSI[−] systems (e.g., LiTFSI-DOL-DME shows ± 1.55 V anodic uncertainty), indicating greater structural inhomogeneity compared to more rigid anions such as PF₆[−].

These anion-specific effects reveal fundamental principles governing electrochemical stability limits. Across

all twelve electrolytes studied, the cathodic limit is set by solvent-separated ion pairs, whereas the anodic limit is governed by contact ion pairs, as shown in Fig. S2. Beyond pure anion effects, the interplay between anion identity and solvent chemistry leads additional complexity in mixed and functionalized systems.

Carbonate mixtures like LiPF_6 in EC-EMC exhibit dramatically different stability windows (5.10 V difference) compared to single-component systems (LiPF_6 -EC: 8.71 V, LiPF_6 -EMC: 9.29 V), with the mixture showing a substantially increased cathodic limit (5.05 V) that reflects cooperative solvation effects. Similarly, fluorine-functionalized electrolytes exhibit systematic shifts in both HOMO and LUMO distributions relative to their parent solvents, with fluorination generally improving cathodic stability but showing variable effects on anodic limits depending on the specific anion-solvent combination. Notably, the model resolves these differences without explicit supervision, capturing the impact of additives on ESWs and demonstrating its ability to distinguish subtle but chemically meaningful variations between single component electrolytes and their mixtures.

Solvation size effects

To evaluate how the size of the solvation environment affects the frontier orbital energies, we analyze variations in predicted HOMO, LUMO, and HOMO-LUMO gap as a function of the solvation shell size and composition (Fig. 3F and G). Li-containing structures were carved out from the MD trajectory to include only the primary solvation shell (coordinated solely by solvent molecules or by one anion), the secondary solvation shell (structures extended to the second-shell distance), and secondary solvation including an anion.

For SSIPs (Fig. 3F(left panel)), the frontier orbital energies exhibit a systematic dependence on solvation shell size. As the environment expands from primary to secondary coordination, the HOMO energy increases substantially, by 2 to 3 eV, indicating reduced oxidative stability. This trend becomes more pronounced when anions enter the secondary solvation shell, despite the absence of direct Li-anion contact. The upward shift in the HOMO reflects long-range electronic coupling mediated through the solvent network, demonstrating that even non-contact ion pairing can strongly perturb the occupied electronic states. In contrast, the LUMO energy in SSIP configurations shows a weaker but still consistent response to the extension of the solvation environment. As shown in Fig. 3F(middle panel), secondary solvation induces only modest shifts in LUMO levels, whereas the inclusion of anions produces a more pronounced effect. Taken together, these effects result in a monotonic reduction of the HOMO-LUMO gap with increasing solvation shell complexity. Similarly trends are also observed for CIPs (Fig. 3G), where the HOMO and LUMO energies show a strong dependence on the solvation environment, particularly on the inclusion of extended solvation shells (second solvation shell) that contain anion species.

Partial desolvation effects near electrode interfaces




Figure 4: Desolvation analysis. (A) HOMO, LUMO, and HOMO-LUMO gap as a function of number of solvents in the primary solvation shell for freely solvated Li-ions. Lines are drawn for the systems most sensitive to desolvation (LiFSI-DEE, LiPF₆-BFE) for highlighting purposes. (B) HOMO, LUMO, and HOMO-LUMO gap as a function of number of solvents in the primary solvation shell for CIPs. Again, lines are drawn for the systems most sensitive to desolvation in the CIP configuration, which include LiFSI-DEE, LiPF₆-BTFE and LiTFSI-DEE.

Having established that electrochemical stability is governed by solvation environments and their associated frontier orbital distributions, we examine how interfacial desolvation process during typical battery operation may further modulate the electrochemical stability of solvated Li species. Structures with different coordination numbers were sampled from our classical MD simulations, representing the range of solvation environments that arise from thermal fluctuations. Fig. 4A and B reveal how the number of solvent molecules in the primary solvation shell influences frontier orbital energies for the freely solvated Li-ions and CIPs, respectively.

For the freely solvated Li-ions (Fig. 4A), partial desolvation induces substantial system-specific shifts in the predicted frontier orbital energies. The systems most sensitive to desolvation are those with ether-based solvents. For example, LiFSI-DEE shows the most dramatic response, with LUMO shift of 2.53 eV (from -4.71 eV at 2 molecules to -2.18 eV at 4 molecules) and gap modulation of 3.23 eV (from 8.20 eV at 1 molecule to 11.43 eV at 4 molecules). Similarly, LiPF₆-BFE exhibits a 1.90 eV HOMO shift from -14.33 eV (2 molecules) to -12.43 eV (4 molecules). In contrast, carbonate-based systems display more modest variations, with LiPF₆-EC and LiPF₆-EMC showing gap changes of only 1.08 eV and 1.18 eV, respectively, over the same desolvation range.

CIPs (Fig. 4B) display similarly strong desolvation effects, with ether-based solvents again exhibiting the greatest sensitivity to coordination number changes. LiFSI-DEE shows the most pronounced LUMO variation of 2.02 eV, shifting from -2.04 eV (1 molecule) to -0.02 eV (4 molecules), with corresponding gap modulation of 1.78 eV (7.71 eV to 9.49 eV). LiPF₆-BTFE demonstrates significant HOMO sensitivity with a 1.60 eV destabilization from -12.06 eV (1 molecule) to -10.46 eV (4 molecules), although its gap value remains relatively constant with only a 0.71 eV variation (11.71 eV to 11.00 eV). In contrast, systems like LiTFSI-DEE show only modest impact from desolvation, with maximal gap variations of 1.33 eV.

These observed differential sensitivities likely arise from variations in local solvation geometry, which lead to greater variations in Li⁺ coordination geometry and, consequently, larger fluctuations in frontier orbital energies. The strong dependence of electronic properties on both solvation number and solvent identity underscores the value of high-throughput screening approaches. In this context, foundation models enable rapid

evaluation of numerous electrolyte candidates across diverse solvation states, and can identify promising systems for subsequent investigation with more computationally demanding yet more accurate physics-based methods. Moreover, the substantial orbital energy variations observed during desolvation highlight that accurate prediction of interfacial electrochemical stability requires explicit modeling of partially desolvated states, suggesting that future predictive simulations would benefit from incorporating electrode proximity effects beyond bulk solvation sampling.

Statistical ensemble perspective and implications for physics-based modeling protocols

The substantial variance in the predicted frontier orbital energies discussed above highlights that electrochemical stability is governed by a statistical ensemble of local solvation environments rather than by a single characteristic molecular structure. The diversity of coordination chemistry and geometry, including Li–O distances, solvent orientation angles, the degree of ion pairing, and the presence or absence of secondary-shell molecules, revealed by extensive sampling of solvation environments in the MD simulations collectively contributes to the spread in predicted orbital energies, albeit with different weights. Critically, the electrochemically relevant states are not the distribution means but the extreme tails: the lowest-lying LUMOs that define the cathodic limit and the highest-lying HOMOs that define the anodic limit. These extreme configurations, which may constitute only a small fraction of the sampled structures, dictate when and where electrochemical decomposition initiates. Consequently, the macroscopic stability window is controlled by rare but thermally accessible solvation environments (SSIPs for reduction and CIPs for oxidation), directly linking microscopic fluctuations in solvation structure to the observed electrochemical limits.

This statistical perspective underscores fundamental limitations of first-principles modeling of electrolyte stability. Static DFT calculations typically operate on energy-minimized clusters and, by construction, cannot access the distribution tails that ultimately determine electrochemical failure, leading to systematic underestimation of both reduction susceptibility (by missing low-lying LUMO configurations) and oxidation susceptibility (by missing high-lying HOMO configurations). The strong dependence of frontier orbital energies on solvation-shell size and ion-pairing structure further indicates that small, isolated molecular clusters are insufficient to capture the electronic environments experienced by solvated species under practical conditions. Accurate prediction of HOMO and LUMO energies therefore requires inclusion of extended solvation shells, explicit representation of ion pairing, and dynamical sampling of coordination environments. Although DFT can, in principle, describe these effects, electronic-structure calculations on large, thermally sampled solvation clusters rapidly become computationally prohibitive. As a result, many existing studies rely on static or undersolvated cluster models that underestimate the impact of collective solvation and speciation. In this regard, ensemble-based ML approaches offer a unique opportunity to overcome these limitations by efficiently sampling the full distribution of solvation configurations that governs electrochemical stability under realistic operating conditions.

Insights for ML model selection

Finally, we emphasize the necessity of explicitly encoding solvation geometry in the ML model by directly comparing our structure-based POS-EGNN model with the SMILES-based SMI-TED model [16] (Fig. 5). The SMI-TED model (SI, Section 9) operates on one-dimensional molecular string representations (SMILES) of the solvent molecules and therefore cannot capture three-dimensional coordination geometry, ion pairing, or solvation-shell structure. This architectural limitation leads to systematic deviations from the POS-EGNN predictions, underscoring the critical role of explicitly encoding structural information in the ML models for accurate prediction of electronic properties.

For HOMO energies, the SMI-TED model exhibits deviations ranging from -5.83 to $+2.14$ eV relative to POS-EGNN across the solvent systems examined. The largest discrepancies occur for G2, with a -5.83 eV deviation (POS-EGNN: -12.64 eV; SMI-TED: -6.81 eV), DOL–DME with -5.50 eV (POS-EGNN: -12.39 eV; SMI-TED: -6.89 eV), and BFE with -4.56 eV (POS-EGNN: -11.80 eV; SMI-TED: -7.24 eV). Carbonate solvents show moderately large errors of -3.21 to -3.51 eV. Only BTFE deviates in the opposite direction ($+2.14$ eV), indicating that the one-dimensional representation fails to capture stabilizing effects associated with specific solvation geometries.

LUMO energy predictions show similarly large discrepancies, spanning -5.41 to $+2.41$ eV. The EC–EMC mixture exhibits the largest deviation at -5.41 eV (POS-EGNN: -4.17 eV; SMI-TED: $+1.24$ eV), followed by DOL–DME at -4.01 eV and G2 at -3.86 eV. Ether systems consistently show deviations exceeding

−2.7 eV (DEE: −2.83 eV; DEE-DPE: −2.85 eV; BFE: −2.71 eV), whereas carbonate solvents exhibit more modest errors of −1.81 to −1.85 eV. Again, BTFE represents an outlier with a +2.41 eV deviation.

Although HOMO–LUMO gap predictions benefit from partial error cancellation, systematic deviations persist, ranging from −2.51 to +2.08 eV. EC–EMC shows the largest underestimation at −2.51 eV (POS-EGNN: 6.52 eV; SMI-TED: 9.03 eV), while EMC exhibits the largest overestimation at +2.08 eV (POS-EGNN: 11.09 eV; SMI-TED: 9.01 eV). Even for systems with relatively small gap errors, such as BTFE (+0.20 eV) and DPE (+0.23 eV), the underlying HOMO and LUMO predictions contain large compensating errors exceeding 2 eV.

The fundamental limitation of SMILES-based models arises from their one-dimensional representation of chemical information. A SMILES string specifies the molecular graph of the solvent but contains no information about how it coordinates the Li ion or how it orients relative to coordinating anions. Yet, it is precisely these geometric details that ultimately determines frontier orbital energies through effects such as orbital hybridization and electrostatic screening.




Figure 5: Comparison of mean predictions obtained from the POSEGNN model and the SMI-TED model

Conclusion

The study introduces a structure-aware ML framework leveraging the POS-EGNN model to predict frontier orbital energies (HOMO, LUMO, and HOMO-LUMO gap) of battery electrolytes. By direct encoding geometric information and incorporating ensembles of extended solvation configurations, our framework bridges the gap between traditional quantum mechanical methods and highly scalable AI-based screening approaches, with enhanced interpretability of ML predictions that allow robust derivation of structure–property relationships for the rational design of advanced battery electrolyte systems.

Through systematic analysis of twelve battery-relevant electrolyte systems spanning carbonates, glymes, ethers, and fluorinated solvents with FSI^- , TFSI^- , and PF_6^- anions, we quantify how local coordination environment and species determines frontier orbital energies. Across all systems examined, it is found that SSIPs control the cathodic limit whereas CIPs determine the anodic limit, with coordination-dependent energy variations spanning 2 to 3 eV. Beyond the primary solvation effect, our analyses also highlight the critical role of extended solvation environment, which can modulate the orbital energies by up to 3 eV through long-range electronic coupling mediated by the connected solvent network. Analysis of partial desolvated states, which naturally occur during battery’s operation at the electrode-electrolyte interfaces, further reveals system-specific sensitivities to coordination number changes, with fluorinated ethers exhibiting particularly strong responses. These results indicate that bulk-phase sampling alone is insufficient to fully capture the electrochemical stability of relevant solvated species and to reliably predict propensities for interfacial degradation.

Our framework addresses a fundamental limitation in both traditional quantum chemistry and previous ML approaches for electrolyte screening. It is shown that the frontier orbital energies are not determined by a single molecular representation of the solvated species, but rather, by the collective interplay of ion-pair interactions, coordination geometry, and the extended solvation environment. Directly encoding these molecular details in the ML model is therefore essential for accurate downstream predictions; correspondingly, ML models based solely on SMILES representations are fundamentally unsuitable for electronic-structure prediction. By training equivariant graph neural networks on the OMol25 database filtered for battery-relevant

chemistries, our model explicitly incorporates geometric features while learning transferable representations of solvation environments from large and diverse datasets. From a practical perspective, this work establishes a scalable pipeline for high-throughput screening of electrochemically promising electrolyte candidates at computational costs that are 3–4 orders of magnitude lower than DFT, while maintaining sub-0.6 eV accuracy in frontier orbital assessment. Nevertheless, we acknowledge several limitations of the current model that warrant further development in future work. First, expanding the training dataset to include a greater diversity of mixed-solvent coordination environments would improve transferability for screening multicomponent electrolytes used in commercial batteries. Second, incorporating temporal information from direct MD simulations of activated desolvation processes may be necessary to capture transient coordination events that govern electrolyte decomposition pathways and interfacial ion-transport kinetics. Third, extending the predictive scope beyond frontier orbital energies to include reaction barriers, decomposition products, and solid–electrolyte interphase properties would enable comprehensive computational screening of electrolyte stability and performance across the full operating window of next-generation batteries.

Conflicts of interest

There are no conflicts to declare.

Acknowledgments

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Portion of the work was supported by Laboratory Directed Research and Development funding under project number 23-SI-002 and portion of the work was sponsored by the U.S. Department of Energy, Office of Electricity (OE), Energy Storage Division. Computing support was provided by the Lawrence Livermore National Laboratory Institutional Computing Grand Challenge program and resources sponsored by the Department of Energy and located at the National Laboratory of the Rockies.

Data availability

The data supporting this article have been included as part of the Supplementary Information.

References

- [1] Xiaoyan Ren et al. “A Biocompatible Deep Eutectic Electrolyte Enables Ultra-Fast Charging in Lithium-Ion Batteries”. In: *Advanced Functional Materials* (2025), p. 2500464.
- [2] Kai Li et al. “Machine learning-guided discovery of ionic polymer electrolytes for lithium metal batteries”. In: *nature communications* 14.1 (2023), p. 2789.
- [3] Jiadong Shen et al. “Machine-Learning-Guided Screening of Advantageous Solvents for Solid Polymer Electrolytes in Lithium Metal Batteries”. In: *Nano Letters* 25.19 (2025), pp. 7801–7809.
- [4] Thuy Duong Pham et al. “Practical High-Voltage Lithium Metal Batteries Enabled by Tuning the Solvation Structure in Weakly Solvating Electrolyte”. In: *Small* 18.14 (2022), p. 2107492.
- [5] Zhao Zheng et al. “Deciphering coulombic efficiency of lithium metal anodes by screening electrolyte properties”. In: *Angewandte Chemie International Edition* 64.30 (2025), e202507387.
- [6] Yanyan Hu et al. “Identification of Potential Electrolyte Additives via Density Functional Theory Analysis”. In: *ChemistrySelect* 8.14 (2023), e202300098.
- [7] Yuxuan Cao et al. “Theoretical Screening for Electronic and Solvation Characteristics of Common Molecules as Electrolyte Additives and Co-Solvents for Alkali Metal Batteries”. In: *Advanced Theory and Simulations* (2025), e00767.

- [8] Xiaozhao Liu et al. “Exploring solvation structure and transport behavior for rational design of advanced electrolytes for next generation of lithium batteries”. In: *Applied Physics Reviews* 11.2 (2024).
- [9] Suyeon Ju et al. “Application of pretrained universal machine-learning interatomic potential for physicochemical simulation of liquid electrolytes in Li-ion batteries”. In: *Digital Discovery* (2025).
- [10] Amil Merchant et al. “Scaling deep learning for materials discovery”. In: *Nature* 624.7990 (2023), pp. 80–85.
- [11] Daniel S Levine et al. “The open molecules 2025 (omol25) dataset, evaluations, and models”. In: *arXiv preprint arXiv:2505.08762* (2025).
- [12] Jerret Ross et al. “Large-scale chemical language representations capture molecular structure and properties”. In: *Nature Machine Intelligence* 4.12 (2022), pp. 1256–1264.
- [13] Xiwang Chang et al. “Integrating Molecular Dynamics and Machine Learning for Solvation-Guided Electrolyte Optimization in Lithium Metal Batteries”. In: *Advanced Science* (2025), e04997.
- [14] Alexius Wadell et al. “Foundation Models for Discovery and Exploration in Chemical Space”. In: *arXiv preprint arXiv:2510.18900* (2025).
- [15] Qidi Wang et al. “Interface chemistry of an amide electrolyte for highly reversible lithium metal batteries”. In: *Nature communications* 11.1 (2020), p. 4188.
- [16] Eduardo Soares et al. “An open-source family of large encoder-decoder foundation models for chemistry”. In: *Communications Chemistry* 8.1 (2025), p. 193.
- [17] IBM. *Position-based Equivariant Graph Neural Network (pos-egnn)*. URL: <https://github.com/ibm/materials>.
- [18] Nikhil Rampal et al. “Structural and transport properties of battery electrolytes at sub-zero temperatures”. In: *Energy & Environmental Science* 17.20 (2024), pp. 7691–7698.
- [19] Seiji Takeda et al. “Foundation model for material science”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 37. 13. 2023, pp. 15376–15383.
- [20] Sarp Aykent and Tian Xia. “GotenNet: Rethinking Efficient 3D Equivariant Graph Neural Networks”. In: *The Thirteenth International Conference on Learning Representations*. 2025. URL: <https://openreview.net/forum?id=5wxCQDt6Mo>.