

An Accurate Charge-Aware Machine-Learning Interatomic Potential for the Reduction of Li-ion Battery Electrolytes in Solution

Yujing Wei^{1,2}, John L. Weber¹, James M. Stevenson¹, Zachary K. Goldsmith², Xiaowei Xie¹, Leif D. Jacobson¹, Richard A. Friesner²

1. Schrodinger (United States)

2. Columbia University

Abstract

Machine learning interatomic potentials (MLIPs), also known as machine learning force fields (MLFFs), offer scalable means of simulating complex systems and processes at \textit{ab initio} level accuracy. One such process is the critical yet still poorly understood formation of the solid electrolyte interphase (SEI) at the anode of a Li-ion battery (LIB) during the first charge cycle, where electrochemical reduction of the electrolyte leads to the generation of decomposition products. MLIPs are uniquely poised to atomistically describe these electrochemical processes, as they are not as affected by the same limitations in bonding and electron transfer as classical force fields. Nonetheless, training MLIPs to run accurate dynamics of a condensed phase with two different oxidation states, such as in electrochemistry, is challenging for many architectures. In this work, we show that by using MPNICE, a message passing MLIP architecture with iterative charge equilibration, we are able to accurately (within 1 kcal/mol) train models along two potential energy surfaces (reduced and unreduced) for LIB-relevant electrolyte systems. Importantly, we demonstrate strategies for sampling and training to examples of anion radicals of these species, which often are not centered on any atom (off-center radicals, or OCRs). We additionally discuss well known limitations of global charge equilibration (Qeq) algorithms in erroneously de-localizing charge, and test methods to alleviate the impact on resulting dynamics. Simulations using these models reveal new insights into electrolyte reduction and considerations for the realistic simulation of electron transfer processes in the condensed phase.

Keywords

Machine Learning Interatomic Potential, Lithium Ion Electrolyte, Redox Chemistry

An Accurate Charge-Aware Machine-Learning Interatomic Potential for the Reduction of Li-ion Battery Electrolytes in Solution

Yujing Wei,^{†,‡} John L. Weber,^{*,‡} James M. Stevenson,[‡] Zachary K. Goldsmith,[†]

Xiaowei Xie,[¶] Leif D. Jacobson,[¶] and Richard A. Friesner^{*,‡,†}

[†]*Columbia University, New York NY 10027, United States*

[‡]*Schrödinger Inc., New York NY 10036, United States*

[¶]*Schrödinger Inc., Portland OR 97204, United States*

E-mail: jack.weber@schrodinger.com; raf8@columbia.edu

Abstract

Machine learning interatomic potentials (MLIPs), also known as machine learning force fields (MLFFs), offer scalable means of simulating complex systems and processes at *ab initio* level accuracy. One such process is the critical yet still poorly understood formation of the solid electrolyte interphase (SEI) at the anode of a Li-ion battery (LIB) during the first charge cycle, where electrochemical reduction of the electrolyte leads to the generation of decomposition products. MLIPs are uniquely poised to atomistically describe these electrochemical processes, as they are not as affected by the same limitations in bonding and electron transfer as classical force fields. Nonetheless, training MLIPs to run accurate dynamics of a condensed phase with two different oxidation states, such as in electrochemistry, is challenging for many architectures. In this work, we show that by using MPNICE, a message passing MLIP architecture with iterative charge equilibration, we are able to accurately (within 1 kcal/mol) train

models along two potential energy surfaces (reduced and unreduced) for LIB-relevant electrolyte systems. Importantly, we demonstrate strategies for sampling and training to examples of anion radicals of these species, which often are not centered on any atom (off-center radicals, or OCRs). We additionally discuss well known limitations of global charge equilibration (Qeq) algorithms in erroneously de-localizing charge, and test methods to alleviate the impact on resulting dynamics. Simulations using these models reveal new insights into electrolyte reduction and considerations for the realistic simulation of electron transfer processes in the condensed phase.

1 Introduction

Machine learning interatomic potentials (MLIPs) are expanding the frontiers of molecular simulations, offering close to *ab initio* accuracy at a fraction of the computational cost.¹⁻⁶ The rapid development of MLIPs has spearheaded simulations that bridge the larger time scales accessible through classical force fields while maintaining the advantages of quantum chemistry, which is crucial for the atomistic understanding of a myriad of applications including catalysis,^{7,8} solution phase dynamics,^{9,10} and drug discovery.^{11,12} In particular, MLIPs have gained substantial momentum for simulating complex systems such as battery electrolytes, where understanding solvation structures, transport properties, and electrochemical reactions at the atomic level is critical for designing safer and more efficient energy storage devices.¹³⁻¹⁶ The formation of the solid electrolyte interphase (SEI) at the anode during the first charge cycle of a Li-ion battery (LIB) is especially important for battery design, yet is still poorly understood and difficult to experimentally characterize.¹⁷ Modeling Faradaic processes, such as SEI formation, will require MLIPs that can account for complex redox chemistry.

While MLIPs without any explicit charge features can effectively train to specific ionic species, the treatment of dynamic and variably charged species such as those relevant for redox chemistry requires a local description of charge and long range electrostatics. A com-

mon strategy used by modern MLIPs, for example fourth generation (4G) HDNNPs^{18,19} and related neural network potentials^{20–26} is combining geometric, atom-centered descriptors, with explicit and equilibrated local atomic charges. This approach also allows the natural inclusion of long-range electrostatic effects. The vast majority of MLIPs which incorporate equilibrated atomic charges make use of the charge equilibration (Qeq)²⁷ method, or some variation thereof.^{20,22}

Accurately modeling the behavior of condensed systems including multiple localized redox states of a given chemical species is particularly challenging for MLIPs. For every configuration, the MLIP must be able to accurately portray two or more distinct potential energy surfaces, and equilibrate charge to effectively obtain the adiabatic ground state. While charge-aware models have been very successful in treating a variety of explicitly charged systems, their application to complex charge transfer events has been limited. Recently, Kocer *et al.* trained a 4G-HDNNP to partition distinct fractional charge to Ferric and Ferrous ions based on the global number of counter ions in an aqueous solution, with preliminary simulations indicating discrete charge transfer events in Fe_2Cl_5 simulations.²⁸

Recently, we introduced MPNICE, an MLIP with explicit and equilibrated atomic charges which additionally incorporates message passing to generalize to arbitrary elements.²¹ MPNICE has been shown to accurately train to predict vertical ionization energies and electron affinities for the ChEMBL-20 set of gas phase drug-like molecules dataset.^{21,22,29} Here we train an MPNICE model for the systematic reduction of LiPF_6 in ethylene carbonate (EC), with the goal of simulating localized reduction of distinct Li-EC complexes and the resulting dynamics in a solvated environment. Similar to previous work,¹⁶ we train on gas phase clusters, now explicitly including the one-electron reduced states. In doing so we encounter metastable reduced states which fail to localize the excess electron on any atomic center, instead populating diffuse orbitals which are not associated with any particular atom, similar to a dipole bound anion or solvated electron.^{30–37} These states present a challenge to effective sampling of the reduced PES, as well as training. As the charge representation of

MPNICE (and to our knowledge all related charge equilibrated MLIPs) is restricted to atom centered charges, it is unable to effectively represent the charge distribution of these “off center radicals” (OCRs). We show that it is possible to address these challenging systems without modifying the atom-centered charge model by identification and sorting of clusters exhibiting OCR character, versus typical “atom centered radicals” (ACRs), during dataset creation.

In this work we find that an MPNICE model trained to a dataset consisting of $\text{Li}_p(\text{EC})_q(\text{PF}_6)_r$ gas phase clusters exhibits near *ab-initio* accuracy on both closed shell and radical potential energy surfaces. We then investigate the application of this model to simulations of a single radical in the condensed phase and find that Qeq can distribute fractional charge throughout the entire electrolyte system, leading to multiple ECs bearing excess charge for a simulation with only one excess electron. This spurious behavior is caused by the well-known tendency of Qeq-based equilibration to lead to fractional charge at infinite separation and motivates further advances in MLIP design. We explore techniques to mitigate this error by partially enforcing the exact constraint of integer charge at infinite separation through data augmentation during training, finding that this can effectively, although imperfectly, alleviate the tendency to over-delocalize excess charge.

This paper is organized as follows: In Section 2, we start in 2.1 with identifying and characterizing OCRs. The spin density, energy, and geometry of these species are analyzed, and the management of these species during dataset construction is described. Section 2.2 presents gas phase cluster benchmarks of the resulting model, MPNICE-RE (Reduced Electrolyte) that is trained to reproduce the energy, dipoles, forces, and charges of reduced electrolyte data. In Section 2.3 we describe spurious charge spreading and outline methods to alleviate this via data augmentation during training. In Section 2.4, we present stable periodic molecular dynamics simulations in the liquid state using the MPNICE-RE model. We discuss errors associated with long-range charge distributions with MLIPs that use Qeq or related approaches, and how data augmentation during training can modify this behavior.

In Section 3, we conclude by summarizing the main findings and discussing future directions. In Section 4, the methods section, we detail the training method and dataset generation, including sampling, filtering, and labeling with DFT. We also discuss the MD and data augmentation methods. This effort is a step towards achieving MLIPs that can describe localized electron transfer and dynamics in solution, an important feature for battery and other electrochemical energy systems.

2 Results and Discussion

2.1 Identification of Off-Center Radicals

In order to effectively simulate a singly reduced electrolyte in the condensed phase, the MLIP must be able to adequately represent the vertical electron affinity of every geometry of a $\text{Li}_x(\text{EC})_y$ cluster. The first step we took to construct a training and test dataset was thus extracting and reducing Li centered clusters from a neutral simulation of LiPF_6 in EC, using a previously trained MLIP shown to be in good agreement with experimental values.³⁸ However, preliminary models exhibited poor performance in reproducing the dipole moment of the reduced clusters, prompting a careful assessment of the system’s redox behavior.

Figure 1A shows the spin density of one example neutral radical $\text{Li}(\text{EC})_4$ cluster, the reduced form of the most common lithium-centered solvation complex composition found in EC electrolytes.^{39–41} The added electron in the reduced cluster resides in a region of space outside of the cluster. The enclosed region of maximal DFT electron density is represented by a centroid that is 6.4 Bohr away from the nearest atom, for the isosurface depicted. The isocontour displayed in figure 1 contains 0.612 electrons, while the largest off-center domain contains 0.609 electrons. Hence, upon a single reduction of EC electrolyte species, the added electron resides in a volume that is incompatible with the standard atom-centered charge description used in charge-aware MLIPs, including MPNICE. The OCR character corresponds to a stable SCF solution and persists with respect to changes in the initial

guess.

We have found that such OCRs persist as the size of the finite cluster extracted is increased, in addition to periodic *ab-initio* single point calculations, ruling out the possibility of this feature being purely an artifact of extracting a gas phase cluster from the condensed phase (see SI Section S1.1). Furthermore, this OCR is present in DLPNO-CCSD(T) calculations, which also excludes the possibility of it being an artifact of DFT or the specific functional used (see SI Section S1.2). This characteristic is shared by the difference density between the $\text{Li}(\text{EC})_4$ and $\text{Li}(\text{EC})_4^+$ state, the SOMO of the reduced state, and the LUMO of the unreduced state (see SI Section S1.2). Therefore, in principle the LUMO of an unreduced cluster can be used to predict the formation of OCR upon reduction, although with some caution, as adding an electron to a small system can be a large perturbation to its orbitals.




Figure 1: Example off-center radical for a cluster of lithium surrounded by 4 EC molecules (Li=magenta, C=grey, O=red, H=white, positive density=violet). Density visualization corresponds to the spin density of the singly reduced state, with a net neutral total charge.

In addition to being present in vertically reduced clusters, OCRs were also seen to persist upon optimization. Reduction of the basis set size, by using $\omega\text{B97X-3c}$,⁴² resulted in a

small fraction of clusters optimizing to ACRs. Figure 2A displays histograms of ω B97X-D3BJ/def2-TZVPD single point energy distributions for the 30 $\text{Li}(\text{EC})_4$ compositions from the holdout test set after optimization with ω B97X-3c. There is a clear energy separation of around 30 mHa (15 kcal/mol) between structures that are optimized to end up as OCR or ACR. This energy gap and inconsistent optimization suggests that OCR are metastable states on the potential energy surface. The energy gap is additionally seen when implicit solvent is used (SI Section S1.3), although it is significantly reduced, presumably due to the stabilization of the large dipole moment of the OCR. As presented in Section 2.4, AIMD simulations suggest that OCR states are nevertheless inaccessible throughout dynamics in the condensed phase. Optimization with ω B97X-3c in implicit solvent resulted in the most consistent localization of the radical on atomic centers (see SI Section S1.3), and was thus used to more effectively sample equilibrium reduced geometries.

As the charge distribution of OCRs are incompatible with MPNICE’s charge representation, the proper identification of these species is critical to curating datasets and charge descriptors. For this work, we sort OCRs and ACRs according to spin density (see Methods), after which dipole labels are excluded in training for the OCRs. This allows sufficient training to reproduce the energetics of the radical PES in these geometries, without effectively contaminating the charge distributions with un-representable data. However, spin densities are relatively information dense and expensive, and screening using this is not always possible or available for large scale datasets. For this particular class of system, we’ve found a strong correlation between geometry and OCR character. Figure 2B shows the improper torsion of each carbonate group for a DFT geometry optimization trajectory for a $\text{Li}(\text{EC})_4$ cluster in gas phase (ω B97X-3c) (the same example as that in the Figure 2A inset). A planar carbonate group has an improper torsion of zero degrees but increases in value as the electron localizes along the optimization trajectory. The EC ring starts puckering and the $\text{C}=\text{O}$ bends away from the plane, corresponding to electron localization in the carbonyl π^* orbital and an increase in sp^3 character. This is consistent with the findings of Balbuena *et*




Figure 2: A) Energy distributions of gas phase ω B97X-D3BJ/def2-TZVPD single points of 30 $\text{Li}(\text{EC})_4$ clusters after DFT (ω B97X-3c) optimization in the gas phase. Inset: starting from the same geometry, example optimization path (single points calculated at ω B97X-D3BJ/def2-TZVPD for each point). The blue/yellow background in the inset indicates OCR/ACR, respectively, at the corresponding cycles in the optimization path. B) Evolution of improper torsion angle of each carbonate for an example geometry optimization trajectory in gas phase ω B97X-3c of one test set cluster. The spin density and geometry (isovalue of $0.001 e^{-1} \text{\AA}^{-3}$) is depicted after certain optimization cycles above the graph. C) Histogram of maximum absolute carbonyl improper torsion for final optimized structures at gas phase ω B97X-3c for the entire 159 molecule test set with diverse $\text{Li}_p(\text{EC})_q(\text{PF}_6)_r$ compositions (see SI Section S2.1 for list). Inset (upper middle) displays the improper dihedral angle being measured, that between the O-C-O plane (yellow) and the O-O-O plane (black) of the carbonyl CO_3 group. The example shown exhibits an improper dihedral angle of 25.9, classified as bent. Inset (lower middle) shows isovalue of $0.015 e^{-1} \text{\AA}^{-3}$ and an example final optimized structure, with the majority of spin density localizing on the “bent” carbonate group.

al. for reduction of a single Li coordinated to EC.⁴³ For simplicity we refer to this geometry as “bent” EC, as opposed to planar. Figure 2C shows a histogram of the maximum improper torsion found in the carbonate groups for the entire 159 molecule test set with various compositions after optimizing with ω B97X-3c (gas phase). The inset shows an example of the improper torsion of a bent carbonate (upper middle) and localization of spin density on said bent carbonate for an optimized cluster (lower middle, large isovalue of $0.015 e^{-1} \text{\AA}^{-3}$ chosen to make the geometry clearer). The ACR geometries exhibit an improper torsion angle of 24-28 degrees, compared to 0-4 degrees for OCR geometries.

2.2 Gas Phase Cluster Benchmarks

Previously, we developed a battery electrolyte dataset for training MLIPs, without the presence of any reduced species.³⁸ After filtering this dataset for improved compatibility with the addition of reduced data (see Methods) and re-training with MPNICE, we obtain the control model MPNICE-E (Electrolyte). The MPNICE-RE (Reduced Electrolyte) is trained using a combination of the electrolyte dataset (460k single point calculations), 40k added reduced data, and 40k added unreduced data (with the same geometries as the 40k added reduced data). The 80k added data is generated through sampling geometry optimization pathways of unreduced clusters, as described above, and AFIR/NEB paths stretching random bonds to ensure adequate representation of equilibrium and out-of-equilibrium OCR and ACR species. The AFIR/NEB bond stretching was found to be crucial for sampling higher energy structures and aid in MD stability (see Section 2.4).

As a test set, we retain a holdout set of 159 clusters extracted from neutral MD using a previous QRNN model,³⁸ ensuring adequate representation of each $\text{Li}_p(\text{EC})_q(\text{PF}_6)_r$ class. This set was then used to generate additional sets, a summary of which can be seen in Figure 3, while the results for MPNICE-E and MPNICE-RE can be seen in Table 1. All cluster benchmark results are for an averaged ensemble of three separate models, trained with different random seeds.

Test sets i. and ii. are equilibrium geometries in the unreduced and reduced state, obtained via optimization of the initial set v. using $\omega\text{B97X-3c}$ and implicit solvent. The 159 clusters were screened from an initial, slightly larger test set to remove remaining OCR character after optimization in the reduced state. Test sets iii. and iv. represent the reduced and unreduced states at the equilibrium geometries of the unreduced (i.) and reduced (ii.) species, respectively, as would be used for a vertical electron affinity. The electron affinity (EA) here is defined as

$$\text{EA} = E_{\text{unreduced}} - E_{\text{reduced}}, \quad (1)$$




Figure 3: Summary of test sets i. to vi, with an example from each taken from the same path. A total of 159 clusters are selected to make the test set v. from which test sets i. to iv. are derived (see SI Section S2.1 for the compositions). The optimization paths that are not selected for the holdout sets are incorporated in the training set (see Methods). Spin densities (isovalue $0.0015 e^{-1} \text{Bohr}^{-3}$) are depicted for the cluster examples for reduced sets. Test set vi. are extracted from small box AIMD simulations (see Methods), which are initiated from a frame equilibrated with QRNN-E (contains OCR).

Table 1: RMSDs for the 159 cluster test set for each model compared to the reference level of theory, ω B97X-D3BJ/def2-TZVPD, reported to 3 s.f. RMSD values for dipole moments are taken as the RMSD of the L2 norms between the MPNICE model dipole and the reference dipole across the test dataset, with regularized percent RMSD given in parentheses. Force RMSDs are taken as the RMSD of the RMSD of the L2 norm of the force difference vectors (between MPNICE and reference force vectors) across all atoms for each molecule. Energy RMSDs are shown for an ensemble average with 3 members for each model.

Set	Model	Energy (kcal/mol)	Dipole (eÅ)	Force (kcal/mol/Å)
i. unreduced (unreduced geometry)	MPNICE-E	1.12	0.024(0.91%)	0.5
	MPNICE-RE	0.437	0.02(1.0%)	0.4
ii. reduced (reduced geometry)	MPNICE-E	70.5	3.01(90.2%)	38.5
	MPNICE-RE	0.688	0.0437(1.2%)	1.0
iii. reduced (unreduced geometry)	MPNICE-E	21.6	4.68(135%)	8.7
	MPNICE-RE	3.26	4.08(111%)	6.7
iv. unreduced (reduced geometry)	MPNICE-E	1.52	0.0464(0.85%)	1.9
	MPNICE-RE	0.547	0.0322(0.71%)	0.8
v. MD (unreduced)	MPNICE-E	0.775	0.0272(0.45%)	0.6
	MPNICE-RE	0.566	0.0262(0.46%)	0.7
vi. AIMD (reduced)	MPNICE-E	115	2.96(115%)	29.4
	MPNICE-RE	2.56	0.144(5.81%)	3.1

i.e. the difference in gas phase electronic energies for a given geometry. Example ORCA input files for optimization and single point energies can be found in the SI Section S5.

MPNICE-RE shows excellent agreement for energies, dipoles and forces for these equilibrium tests (see Table 1). As expected, MPNICE-E obtains good results for the unreduced species. Interestingly, MPNICE-RE, despite being trained on multiple charge states, improves significantly on the results for the unreduced clusters, which we attribute to the slightly expanded representation of off-equilibrium geometries. For the reduced sets (ii. and iii.), MPNICE-RE represents a significant improvement. For test set ii. which includes no OCRs, MPNICE-RE achieves an energy RMSD of 0.69 kcal/mol, and dipole/force RMSDs on par with, although slightly higher than, the unreduced clusters. In contrast, test set iii. includes all OCR structures, resulting in significantly higher errors for the energy and force, with almost no improvement on the dipole moment versus MPNICE-E.

Figure 4 displays parity plots for the adiabatic EA and vertical EA (where the geometry

is the optimized to the reduced state in Figure 4B, i.e. test set ii., iv. and the unreduced state in Figure 4C, i.e. test set i., iii. for MPNICE-RE. Although the errors for the OCR structures are higher, the model is still adequately able to identify that OCRs are higher in energy than ACRs. As we expect these to be metastable states with very low frequency in MD, the decrease in accuracy in total energy for these should result in relatively little distortion of the dynamics. Adiabatic electron affinities (i., ii.) are closely related to reduction potentials, making MPNICE-RE’s precision highly valuable for electrochemical applications, while vertical EAs demonstrate the ability of MPNICE to describe the energetics of two adiabatic charge transfer states at either equilibrium geometry. The MAD of 0.51 kcal/mol for this validation set implies an average model error of less than 0.02 V in reduction potential estimations versus DFT.




Figure 4: Correlation plot for gas phase EA, for the MPNICE-RE model vs DFT at ω B97X-D3BJ/def2-TZVPD. A) Adiabatic, B) vertical at the geometry optimized in the reduced state, C) vertical at the geometry optimized in the unreduced state (off-center radical once reduced). Note that the magnitude of EA is closely correlated to the cluster composition; for example, clusters with composition $\text{Li}_1(\text{PF}_6)_2(\text{EC})_x$ exhibit the lowest EAs due to the proximity to counterions. Errors in EA, however, were not found to correlate significantly with composition, and the residuals can be seen in Section S2.3

Test set v. are MD geometries extracted from the original electrolyte model run with QRNN, which can be referred to as QRNN-E,³⁸ and are expected to be close to the equilibrium geometries for the unreduced state as the QRNN-E is trained to energies, dipoles and forces for unreduced species only. These geometries are also used as the starting point in optimization to obtain test sets i. and ii. MPNICE-E and MPNICE-RE obtain similar

errors for this set and are similar to that of set i. As an additional test set (vi.) for reduced clusters away from equilibrium but still accessible in MD, we extract clusters from small-box AIMD runs with PBE-D3^{44,45} and a plane wave basis set (see Methods for details). PBE has been shown to have a significantly reduced barrier for ring opening^{46,47} compared to hybrid functionals, and thus this set additionally contains some reacted EC. Although we do include some AFIR/NEB paths stretching in the training data, this was primarily to enhance MD stability, and so the higher RMSD of MPNICE-RE for set vi. is primarily attributed to reacted clusters, albeit still good with a 2.56 kcal/mol RMSD.

Finally, we note that MPNICE-RE is able to optimize in the reduced state to ACR states, starting from the unreduced equilibrium geometry. In this case, we optimize, starting from test set v., and 51% of the structures after MPNICE-RE optimization are no longer OCR, with the characteristic energy gap that accompanies overcoming the plateau in energy between OCR and ACR. The analogy of Figure 2A for MPNICE-RE optimization of the test set can be found in the SI Section S1.4. For reference, optimization with ω B97X-3c in gas phase and implicit solvent, returns 17%, and 92% ACR for the same subset of the training set (SI Section S1.3), respectively. ω B97X-D3BJ/def2-TZVPD optimization was not tested on the test set, as preliminary tests on the training set (SI Section S1.3) returned $< 3\%$ ACR after optimization in implicit solvent.

2.3 Size Consistency and Charge Spreading

Qeq²⁷ is used for global charge distribution in many modern MLIPs, including MPNICE, i.e. it allows for charge description beyond the MLIP cutoff radius, which is crucial for modeling properties such as polarization, charge transfer and intermolecular interactions. However, it is well known that despite its simplicity, global charge conservation, low computational cost, and compatibility with MLIP architectures, the electronegativity equalization principle that Qeq uses permits unphysical charge transfer across arbitrary distances.^{27,48–52} Qeq can be described as “metallic”, as charge redistributes across the entire system, regardless of bonding




Figure 5: Scheme for ISA-I and ISA-II. The net charge is restricted for the total such that only one of the molecules is reduced. For this example shown above (two $\text{Li}(\text{EC})_4$ fragments), the net charge on the total system is 1, and the charge assignment during training is such that the cluster with the more favorable electron affinity is assigned the reduced (e.g. 0) charge.

or distance. Even for small systems, fractional charge from Qeq can result in incorrect bond dissociation energies at infinite separation, and, while carrying out simulations of larger systems, this non-size consistent behavior of partial charges affects energies, dynamics and reactivity. For systems with consistent charge, such as the unreduced $\text{Li}(\text{PF})_6$ in EC electrolyte, this fractional charge is consistent and has not resulted in perturbed dynamics in the condensed phase, with QRNN models obtaining bulk properties for electrolytes in good agreement with experiment.³⁸

In initial simulations with a single excess electron, MPNICE-RE resulted in non-physical delocalization of fractional charge on multiple ECs (see Section 2.4). To investigate this, we tested for size consistency and fractional charge spreading on singly reduced combinations of clusters from holdout test sets i. and ii. separated by 1000 Å (Figure 5). The tests were separated into two classes of cluster combinations; bent-planar, for which one cluster has a

bent EC, which should localize the excess electron on the bent cluster, and bent-bent, for which the localization of the electron depends on the relative EAs of the two clusters. As seen in Table 2, MPNICE-RE exhibits significant size consistency error, $E_{AB} - (E_A + E_B)$, for both combinations of clusters tested. The bent-bent set has a much larger error in integer charge, effectively splitting the charge between the two clusters, rather than localizing on the one which has a greater electron affinity.

Many methods have been proposed which effectively correct Qeq by reformulating it in terms of local charge transfer between atomic sites; however, such methods rely on an auxiliary set of initial formal charges, which already satisfy asymptotic constraints, and represent discontinuities when treating reactions or charge transfer.^{48,53} In lieu of explicitly setting reference charges, we attempt to alleviate this behavior by augmenting training with asymptotic constraints in the loss function. This is similar in concept to enforcing symmetry constraints via data augmentation.^{54,55} We refer to this as infinite separation data augmentation (ISA), and fine-tuned MPNICE-RE using two implementations, corresponding to the two classes of cluster in Figure 5. We refer to these as ISA-I and ISA-II, which involve different criteria for eligible data to be paired to form the augmented datapoints, with ISA-I having less diversity. ISA-I only pairs bent (i.e. clusters with a bent EC) and planar (clusters with no bent EC present) data and assigns the radical charge to the member of the pair containing the bent EC, while ISA-II also contains bent-bent pairs. See Methods for details. The datasets for these models contain the original MPNICE-RE dataset, and in addition 20% augmented data consisting of infinitely separated pairs of the original data points (A and B) where size consistency,

$$E_{AB} = E_A + E_B, \tag{2}$$

is encouraged through the augmented label and loss function.

As seen in Table 2, ISA training improves significantly on the size consistency and fractional charge of MPNICE-RE. Test sets I. and II. consist of generated pairs from holdout test sets i. and ii., and according to the ISA-I criteria and ISA-II criteria that does not

overlap with ISA-I, respectively—i.e. Bent-Planar pairs and Bent-Bent pairs. See SI Section S2.2 for an example. The II. test, and ISA-II training (bent-bent) can be unforgiving, as the relative electron affinities between two bent clusters can be extremely similar, unlike the ISA-I method where the EA between bent and non-bent clusters are more distinct. In training to bent-bent ISA, MPNICE-RE-ISA-II obtains a reasonable 3.53 kcal/mol size consistency error, but sacrifices some performance for the I. test (Bent-Planar) compared to MPNICE-RE-ISA-I, no longer obtaining < 1 kcal/mol size consistency error for bent-planar ISA.

ISA places the responsibility for size consistency correction on the neural network, and is explicitly dependent on the iterative nature of charge equilibration used in MPNICE, which facilitates implicit learning from both local and global environmental influences required for this correction. If the architecture were not to include input or intermediate parameters explicitly containing both global (Qeq) information and local geometric information, augmentation like this would likely result in conflicts with the original dataset, biasing the results on separated clusters. Table 3 summarizes the results for the MPNICE-RE-ISA-I and MPNICE-RE-ISA-II models on the original test sets in table 1. Although the RMSDs are in some cases slightly increased with respect to MPNICE-RE, the trends match closely, and all total energies are reproduced well within 1 kcal/mol, aside from sets iii. and vi. which contain OCRs and reactive intermediates respectively, as discussed above.

Table 2: RMSDs for each model for the two size consistency augmented test sets to 3 s.f. Fractional charge refers to the absolute deviation of both species in the pair from integer total charge; for example, the total molecular charge for a pair of infinitely separated molecules are $n + \delta$ and $m - \delta$, where n and m are integers and δ is the fractional charge. See SI Section S2.2 for an example.

Set	Model	Size Consistency Error (kcal/mol)	Fractional Charge
I. Bent-Planar	MPNICE-E	39.7	0.402
	MPNICE-RE	10.6	0.079
	MPNICE-RE-ISA-I	0.759	0.00239
	MPNICE-RE-ISA-II	2.31	0.0123
II. Bent-Bent	MPNICE-E	16.9	0.399
	MPNICE-RE	26.8	0.315
	MPNICE-RE-ISA-I	35.5	0.406
	MPNICE-RE-ISA-II	3.53	0.101

Table 3: RMSDs for the 159 cluster test set for each model compared to the reference level of theory, ω B97X-D3BJ/def2-TZVPD, reported to 3 s.f. The MPNICE-RE model has 540k datapoints, curated by adding 40k datapoints with 2 oxidation states (reduced and unreduced) each to the control MPNICE-E model of 460k datapoints (unreduced only). MPNICE-RE-ISA-I and MPNICE-RE-ISA-II have 648k datapoints, from augmenting 20% of MPNICE-RE data and adding to the dataset. RMSD values for dipole moments are taken as the RMSD of the L2 norms between the MPNICE model dipole and the reference dipole across the test dataset. Force RMSDs are taken as the RMSD of the L2 norm of the force difference vectors (between MPNICE and reference force vectors) across all atoms for each molecule. RMSDs are shown for an ensemble average with 3 members for each model. Vertical EA refers to energy differences between two electronic states (reduced and unreduced) for each species and hence only energies are reported.

Set	Model	Energy (kcal/mol)	Dipole (eÅ)	Force (Ha/Å)
i. unreduced (unreduced geometry)	MPNICE-RE-ISA-I	0.509	0.0228	0.000761
	MPNICE-RE-ISA-II	0.54	0.0199	0.000851
ii. reduced (reduced geometry)	MPNICE-RE-ISA-I	0.678	0.0451	0.0019
	MPNICE-RE-ISA-II	0.806	0.0414	0.00214
iii. reduced (unreduced geometry)	MPNICE-RE-ISA-I	3.47	4.16	0.0114
	MPNICE-RE-ISA-II	3.62	4.24	0.0113
iv. unreduced (reduced geometry)	MPNICE-RE-ISA-I	0.587	0.0382	0.00147
	MPNICE-RE-ISA-II	0.651	0.0305	0.00287
v. MD (unreduced)	MPNICE-RE-ISA-I	0.643	0.05	0.00126
	MPNICE-RE-ISA-II	0.804	0.0268	0.0015
vi. AIMD (reduced)	MPNICE-RE-ISA-I	2.39	0.13	0.00564
	MPNICE-RE-ISA-II	4.56	0.132	0.00673

2.4 Molecular Dynamics

For efficiency, MD is run with a single MPNICE model, rather than an ensemble as for the cluster benchmark. The MPNICE-RE reduced model is stable in MD for simulations both in the unreduced and singly reduced states, for the duration of every MD run (≥ 1 ns). For both AIMD with PBE-D3 and MPNICE-RE MD, we observe the process of the reduction of EC manifesting in one EC becoming bent after an electron is injected in the simulation. This geometry change is characteristically the same as that described in Figure 2. Figure 6A shows the improper torsion angle vs propagation time in AIMD for in the singly reduced state (simulation box with net charge of -1 and a uniform backing charge). Sample frames from this AIMD run are shown in Figure 6B, a few femtoseconds before and a few femtoseconds after the switch of the bent EC (transfer of electron) as the environment changes. Inclusion of an excess charge via the exclusion of an PF_6 counterion led to similar results - see SI Section S3.

In MPNICE-RE MD, although the localization of the excess electron and corresponding EC geometry change occurs later (a few hundred picoseconds instead of a few femtoseconds, see SI Sections S3.1 and S3.3), a second EC is additionally bent (Figure 6C) over time. This is in contrast to AIMD, where typically only one EC is bent at a given time. For both bent ECs, the MPNICE-RE partial charges on the carbonate carbon are shifted towards more negative values. The sum of charges on each bent EC is -0.5 (see Figure 6C). This sum more so reflects the environments of the bent ECs than the total number of bent ECs, as the additional bent EC did not seem to accept negative charge from the first EC to bend. The other solvent molecules are uniformly slightly more positively charged to accommodate the bent ECs (SI Section S3.1), characteristic of Qeq charge spreading.⁵³ There is also some contribution from Li becoming slightly more positive, but not by an amount that offsets the local -0.5 charge accompanying a second EC bend.

In order to analyze the impact of infinite separation data augmentation on the dynamics of a singly reduced electrolyte, we ran MD with one excess electron starting from six distinct




Figure 6: A) Improper torsion vs time for AIMD with a periodic box of Li₁(EC)₁₄(PF₆)₁, initialized without reduced/bent ECs. Each line represents one carbonate group. At 15.4 ps, the carbonate that is bent switches, representing a discrete electron transfer event. B) Two superimposed frames, of the simulation from A), one a few femtoseconds before the switch (blue) and the other a few femtoseconds afterwards; the carbons involved in the switch are circled in orange. C) Frame from MPNICE-RE MD after two EC (circled in orange) both become bent, with a box of Li₅(EC)₆₃(PF₆)₄. The color scheme shows the MPNICE atom centered partial charges on each atom. Outset zoom-in of one of the bent ECs shows that the MPNICE partial charges on the bent carbonyl C are much more negative compared to other carbons. D) Percentage of time spent with 0, 1, and greater than one bent EC for simulations initialized with zero bent EC. E) Statistics for simulations that start with one bent carbonate. F) Statistics for trajectories starting from frames with multiple bent carbonates.

frames; two without any bent EC, one with 1 bent EC, one with 2 bent EC, one with 3, and one with 5. Qualitatively, the correct behavior is to bend a single EC, localizing the excess electron; we track the number of bent EC by monitoring the improper torsion of each EC over time, with a magnitude $> 11^\circ$ interpreted as bent. The percentages of simulation time spent in each qualitative state (zero, one, or more than one EC bent) for each model class, MPNICE-RE, MPNICE-RE-ISA-I, and MPNICE-RE-ISA-II, for the corresponding three classes of starting frame, can be seen in Figures 6D-F. We run these using three models trained with distinct random seeds for each class of training. Improper torsion angles for each trajectory can be seen in the SI Section S3.3. We additionally tested a control where MD is run without an excess electron, and for all the MPNICE models the control is stable with 0% of frames containing a bent EC.

When starting with zero bent EC (Figure 6D), MPNICE-RE is seen to successfully localize and bend an EC, but additionally results in unphysical states with multiple bent EC. Interestingly, the ISA models are unable to initiate localization given an initial state. While this is in contrast to AIMD, it is expected given the inability of the model to accurately describe the OCR state which is present when there are no bent ECs. In this case, the excess charge is distributed evenly throughout the atoms in the simulation, resulting in minimally perturbed dynamics compared to the unreduced state. Note that the slight indication of bending EC in the ISA-II model is due to short lived fluctuations just beyond the chosen 11° cutoff which do not persist beyond a single frame.

When initialized with one bent EC (Figure 6E), the ISA models effectively stabilize the singly bent state, with 100% of the MPNICE-RE-ISA-I trajectories retaining one bent EC, and two out of three MPNICE-RE-ISA-II models retaining one bent EC. One of three ISA-II models bent two additional EC around 900 ps into the simulation. However, all three bent EC were associated with a single lithium cluster centered on a complex of three Li ions, rather than the previously observed reduction of distinct clusters.

When starting from frames with more than one bent EC (Figure 6F), the ISA models

are able to unbend some EC. In the case of MPNICE-RE-ISA-I, one model (random seed 3) effectively unbent all but one EC within the first 200 ps, thereafter remaining stable for the duration of the simulation. The other two MPNICE-RE-ISA-I models retained multiple bent ECs for the duration of the simulation. For the MPNICE-RE-ISA-II models, one model (random seed 2) consistently unbent all but one EC within the first few frames, retaining one bent EC for the duration of the remaining simulation. The two remaining MPNICE-RE-ISA-II models, those that bent multiple EC when starting with one bent EC (Figure 6E) sometimes resulted in states with zero bent EC, as well as one case with multiple bent ECs, again associated with a single complex of three Li ions. One MPNICE-RE-ISA-II simulation additionally crashed after about 80 ps due to uncontrolled reactivity.

3 Conclusion

We have developed an MLIP model that accurately describes Li ion battery electrolyte species in two oxidation states: the unreduced state and a one-electron reduced state, representing a key step towards describing SEI formation. We demonstrate that this model, MPNICE-RE, which uses the MPNICE architecture, is able to obtain total energies within 1 kcal/mol RMSD compared to the reference level of theory (ω B97X-D3BJ/def2-TZVPD) for a holdout set of gas phase clusters, and also achieves good performance for a test set with reacted clusters. This includes the prediction of both vertical and adiabatic electron affinity, with RMSD of 0.7 kcal/mol for adiabatic electron affinities of the holdout set and $R^2 > 0.999$ parity. Given this performance, this model is a starting point for calculating redox potentials, given the correlation between gas phase redox potentials and electron affinities. Furthermore, MPNICE-RE is able to optimize structures in the reduced state from a starting point which is the equilibrium geometry of the unreduced state, achieving similar results to DFT. Hence, the dipoles, energies and forces are successfully trained to for two oxidation states at the same time.

In developing the training dataset for MPNICE-RE, we determined that meticulous curation and sorting is required to be compatible with the charge description of the model. We discovered radicals where the spin density of the radical electron can not be described by atom-centered charges. These OCRs are metastable, as we discover through geometry optimizations, and the energy minima for OCRs are around 10-15 kcal/mol higher than radicals where the spin density is localized. These OCRs can also be identified geometrically; ACRs where the spin density is localized on one EC site contain a characteristic pyramidal distortion of the carbonyl group which exhibits sp^3 character. Therefore, reduced species at the equilibrium geometries of the unreduced state are quite often OCRs. Based on these findings, we remove dipole and force labels from OCRs during training to remove noise due to the charge description. It’s important to note that such OCRs are not likely to impact the training of models which lack localized atomic charges, such as those which embed the total charge as a global feature.⁶ Note that this does not imply that such models are free from size extensivity errors or erroneous charge delocalization behavior.

MPNICE-RE exhibits stable molecular dynamics (MD) trajectories in multiple redox states of electrolyte media. We found that in contrast to AIMD, MPNICE-RE may localize the excess electron on more than one EC, associated with distinct Li-ion centers, resulting in the characteristic bend on both ECs. This is attributed to the global charge equilibration used in MPNICE, based on Qeq, which is known to over-delocalize charge, and verified by explicitly testing for size extensivity and fractional charge for infinitely separated clusters. We demonstrated that the addition of infinitely separated data in the training set can effectively train MPNICE models to favor integer charges on separate solvation complexes (MPNICE-RE-ISA-I), and even to localize integer charges on the most energetically favorable cluster (MPNICE-RE-ISA-II). These ISA models additionally have a much lower rate of erroneously localizing charge on multiple ECs during condensed phase molecular dynamics.

The ISA models are unable to readily initialize localization when starting from the unreduced state; although this is not ideal, this is arguably somewhat consistent with DFT,

where optimization at the ω B97X-D3BJ/def2-TZVPD level does not necessarily localize the electron. Additionally, among ensembles of three MPNICE models trained using the ISA I and II protocols, the degree of localization on separate EC were mixed. The most promising protocol was ISA II, where one of the three MPNICE-RE-ISA-II models stably localized the single excess electron on one EC regardless of the number of bent EC in the starting frame, as long as at least one was present in the starting frame. These discrepancies in qualitative performance between models trained with the same protocol suggest that the training protocol is still not converged, which can perhaps be mitigated through increased model capacity or dataset curation.

While we have demonstrated that tracking the number of bent EC in a molecular dynamics simulation can be effectively used to test training protocols for charge aware MLIPs, it is by no means a perfect test, inherently relying on assumptions that an excess charge can only localize and bend one EC. With increasing evidence of multi-lithium clustering in EC based electrolytes,⁴⁷ we have included multi-lithium clusters in our training. It is entirely possible that delocalization of a single electron between multiple ECs associated with the same multi-Li cluster, as seen in some MPNICE-RE-ISA-II MD simulations, is physically consistent, and further study is warranted. Enforcing charge localization in the condensed phase through data augmentation during training provides a promising, though imperfect, solution to non-physical charge transfer in MLIP based simulations of redox processes. In the future, more rigorous charge constraints may be required, or investigation into other charge equilibration methods such as QTPIE,⁴⁸ ACKS2,⁵⁶ or other partial charge schemes such as maximally localized Wannier Centers.²⁴

The development of an MLIP for the simultaneous treatment of multiple oxidation states in the condensed phase, and associated electron transfer events, is a critical and challenging problem. MPNICE-RE and the derived models are a starting point for increasingly complex simulations of electron transfer in realistic simulations of Li-ion batteries. While some reactivity has been incorporated through AFIR/NEB to generate the training set, the tran-

sition states from these paths have not yet been validated against DFT, and a fully reactive simulation of electrolyte decomposition is reserved as a subject of future work.

4 Methods

4.1 Density Functional Theory

DFT is run with the ORCA⁵⁷ quantum chemistry package, version 6.0.1 and version 5.0.4 (for the latter, a stability check is run for UKS solutions). The ω B97X-D3BJ DFT functional^{45,58,59} was used for all test set and training set data, combined with the def2-TZVPD basis set.^{60,61} While the original unreduced electrolyte dataset for MPNICE-E and QRNN-E³⁸ is run with Psi4, we observe that ORCA produces more variationally correct SCF solutions for reduced (radical) species (see SI S4.1). We have provided ORCA input files for geometry optimization and single point energy with our chosen settings in the SI Section S5. In addition, we find that ORCA 6 has a higher rate of stable SCF solutions, especially compared to ORCA 5 (see SI Section S4.1), and as stability check can increase the cost of a calculation by up to an order of magnitude, we have not run the stability check while using ORCA 6.

4.2 Filtering of Off-Center Radicals

OCRs are identified using a heuristic where the centroid of the region with the largest integrated electron density (of the DFT spin density) for a chosen isovalue for the clusters of the dataset is more than 1.5 Bohr from the center of the nearest atom. For the reduced dataset, this chosen value is $0.0012 \text{ } e^{-1} \text{ Bohr}^{-3}$. For reference, the isosurface value for the display of spin density chosen in Figure 1 is $0.005 \text{ } e^{-1} \text{ Bohr}^{-3}$. In the cases where the total electron density from this isosurface is less than 0.1 (out of a total 1 electron in the case of OCRs), the isovalue is halved until this criteria is met.

4.3 Geometry Optimization

We use ORCA⁵⁷ to optimize with DFT. The training and test set are optimized with custom tolerance thresholds. `TolE`, `TolRMSG`, `TolMaxG`, and `TolMaxD` are all 10 times looser than the defaults in Orca 6.0, and `StrongSCF` is used for SCF convergence — see SI Section S5 for an example input file for single point energy (used in labeling), compared to SI Section S5.2 for an example input file for geometry optimization with the chosen thresholds. We have tested these thresholds against the defaults with no qualitative changes to off-center radical behavior (which is consistent with the observation that OCRs occur early in the optimization trajectory—see insets in Figure 2) and minimal (within a kcal/mol) changes to final single point energies. The number of maximum optimization cycles set is the larger of 50 and $3N$, where N is the number of atoms, which is the default in ORCA 6.0, while the number of maximum SCF cycles is set to 500.

MLIP optimization is performed with ASE using BFGS algorithm with line search, and a maximum of 1000 optimization cycles, with a convergence force threshold of 0.01 eV/Å.

4.4 Model Training

All MPNICE models use the MPNICE architecture as described in Ref. 21 without modification. MPNICE models were trained using multi task loss function.^{62,63} The loss function is minimized with the AdamaxW optimizer, with a weight decay of 1.0e-3. The batch size used is 64. The number of radial and angular message features used are 16 each. The number of iteration blocks is 3, the node feature dimension is 128, and the MLP dimension is (913, 144, 112, 96, 129), and weight normalization is used.

Training datasets were split into training and validation sets using a random split with a 9:1 ratio. Each ensemble uses a different random seed for train:validation split, and each model shown in this work is trained to 3 independent ensembles (molecular dynamics are run separately for each ensemble member). Test set errors are reported as ensemble averages.

We train to fitted “atomization energy” labels, which results from a linear regression to

normalize and center the energies around zero. We note that we use the QRNN-E dataset³⁸ for the atomic energy fitting and use these same atomic energies for every model. Non-ISA models are trained for 400 epochs. Models are trained to energy, dipoles, and forces (gradients), against the ω B97X-D3BJ functional with a def2-TZVPD basis set. Dipoles and forces are excluded from training for datapoints determined to contain OCRs.

4.5 Training Set Generation

The training set is produced by combining the original battery electrolyte dataset³⁸ (460k datapoints) with the described singly reduced data (40k). We additionally add the same 40k generated data for reduction but in the unreduced state to prevent the association of geometry with charge state. In total, the MPNICE-RE model is trained to a dataset with $\simeq 540$ k datapoints. Training set species with i) stray atoms, ii) unphysical bond lengths (based on covalent radii), and iii) with fragment separation $> 5.2 \text{ \AA}$ are filtered out (see Section 4.7). The reduced dataset consists of the following data, beginning with single and multi Li-centered $\text{Li}_p\text{EC}_q(\text{PF}_6)_r$ clusters extracted from MD run with the unreduced MPNICE-E model (these starting structures are chosen to have a diverse energy spread):

1. 10k from optimization trajectories (with the reduced charge and multiplicity) starting with the unreduced MD geometries. We select points that are > 1 RMSD (structural difference) apart along each trajectory. This will choose more structures early in the optimization trajectory, so we also include the final structure.
2. 30k from AFIR/NEB paths (run with an intermediate version of MPNICE-RE before finalizing the dataset), which are also run with the reduced charge and multiplicity starting (reactant) from the unreduced MD geometries. We use AFIR to break one random bond in the cluster and sample the entire NEB trajectory of 15 frames each. Note that the latter half of this trajectory will be filtered out by the stray atom threshold if it’s a dissociation reaction. The bonds broken are C-H, C-O, Li-O, C-C,

and F-P, of which C-O results in ring-opening.

Cluster extraction follows the protocol of identifying Li centers (specified by a Li-Li cutoff distance to determine whether two or more Li are part of the same cluster), and building a graph such that molecules which contain atoms within a specified cutoff distance from any of the central Li atoms are considered as part of the cluster. The Li-Li cutoff distance used in this work is 4.5 Å and the Li-(other atom) cutoff-distance ranges from 2.5 Å to 4.5 Å for training and testing.

4.6 Test Set Generation

The holdout set consists of 160 starting geometries which were extracted from MD run using QRNN-E—these geometries comprise the test set v. The selection of the holdout set is such that each composition found in the training set is represented in the test set, weighted by frequency and selected randomly. The composition spread of the test set can be found in the SI Section S2.1.

Using these starting geometries, optimization was performed in implicit solvent (using SMD parameters for pure EC⁶⁴) with ω B97X-3c, in both the unreduced and reduced oxidation states for each geometry to form the i. and ii. test sets respectively. This choice was made because this level of theory gives the highest ratio of ACR to OCR after optimization (see SI Section S1.3), compared to ω B97X-D3BJ/def2-TZVPD in implicit solvent and ω B97X-3c in gas phase. The test set was filtered to contain clusters that consist of no OCRs after optimizing with ω B97X-3c in implicit solvent. One geometry dissociated during optimization and was additionally filtered out (see the Dataset Filtering subsection), leaving 159 datapoints in each test set in Table 1. Vertical reduction of test set i. and oxidation of test set ii. result in the test sets iii. and iv, respectively.

We hold out these molecules for training, which means we do not add geometry optimization or NEB trajectories starting from the test set structures to training. Although the repetitive composition of Li-EC-based clusters means that it is not possible to completely

avoid the potential space of overlap with the training set, we generally observe large conformational degrees of freedom and local minima in the training and test dataset. We also ensure that no two holdout set structures are less than 1 Å RMSD different in structure.

Furthermore, we keep structures extracted from small-box AIMD (test set vi.) as a test set, which includes reacted structures as well as under-coordinated structures to show the out of equilibrium distribution test.

4.7 Dataset Filtering

All datasets are filtered for:

- i) Stray atoms: we filter out any datapoint that contains any hydrogen atom that is more than 2 Å separated from the closest non hydrogen atom (i.e. we also filter out stray H₂ molecules), and/or contains any heavy (non hydrogen) atom that is more than 3.2 Å separated from the nearest atom. This filter is an automated detection for dissociated species in the training and test set.
- ii) Unphysical bond lengths: if the ratio of a bond distance divided by the sum of covalent radii is greater than 1.5, or if this ratio is smaller than 0.85, the bond is considered to be unphysical and the structure is removed. Note that this excludes structures sampled from AFIR/NEB with deliberate bond breaking.
- iii) “Infinite” separation: if there is more than one molecule in the datapoint (where molecule is a connected series of atoms according to physical bond lengths as discussed above), the datapoint is removed if any two molecules are separated by more than 5 Å. It is known for the Qeq method (which is used in MPNICE and many other MLIPs) to nonphysically spread charge, and unless special labels are given for enforcing integer charges on uncoupled species, such examples of uncoupled species are removed.
- iv) Removed radicals and Li metal clusters from the MPNICE-E dataset. This includes

the control model shown in the results. One of the main results of this paper is that some difficult species such as radicals requires more careful sampling and filtering.

- v) Force and energy cutoff (only accept lower than) maximum force of 0.1 Ha/Å for any atom in the molecule and 0.01 Ha atomization energy normalized by number of atoms.

4.8 Molecular Dynamics

MLIP MD is run with ASE using an NPT ensemble at 313 K. The starting boxes are generated with disordered system builder (tangled chain) in the Schrödinger Suite, then equilibrated with MPNICE-E in the unreduced state. Simulations (both MLIP and AIMD) with a non-zero net charge are made to be neutral through a uniform backing charge.

AIMD is run with Quantum Espresso (QE)^{65,66} with PBE-D3 and plane wave basis (wavefunction cutoff energy 40 Ry, charge density cutoff energy 320 Ry, ultra-soft pseudopotentials for all elements) with NVT ensemble, controlled by velocity rescaling, and used gamma point only sampling. Gaussian smearing with a broadening of 0.01 Ry was used for electronic occupations. The electronic self-consistency was converged to 1×10^{-5} Ry.

LSDA spin polarization was used and boxes were run with 1 additional negative charge and the default jellium counter charge. The following box compositions were run for sampling and at both 313 K and 600 K (a total of 6 different AIMD runs), for around 30 ps each with a timestep of 1 fs and ionic equations of motion integrated using the Verlet algorithm.

- Li(EC)₁₀(PF)₆
- Li(EC)₁₄(PF)₆
- Li(EC)₁₅ (neutral box)

The starting frames are obtained from NPT equilibration for 100-1000 ps using the unreduced MPNICE-E model.

4.9 Infinite Separation Augmentation (ISA)

We train 2 different models, which we call MPNICE-RE-ISA-I and MPNICE-RE-ISA-II. For both ISA models, the dataset is comprised of the 540k datapoints used in the MPNICE-RE model. 20% of this dataset is augmented and added back to the dataset resulting in a total of 648k training datapoints after augmentation (of which we split 9:1 into train and cross validation, the same as for the non ISA models). The datapoints that are eligible for augmentation are only datapoints that are extracted from MD or from optimization/NEB paths starting from these datapoints, i.e. clusters of molecules consisting of Li, EC, and PF_6 . Most of these datapoints are from the 40k additional geometries (compared to the original electrolyte dataset used in MPNICE-E), however, around at least 10k from the original electrolyte dataset are also applicable. These eligible datapoints are split into two subsets, the bent and planar subsets, depending on if a bent EC (non-zero improper torsion carbonate group) is present in the cluster. In practice, spin densities are computed for all relevant datapoints to separate OCRs and ACRs during training, and this also indicates whether there exists a bent EC (see Figure 1 and 2).

The added augmented datapoints consist of pairs of original datapoints MPNICE-RE into a new datapoint, separated by a distance 1000 Å (an arbitrarily chosen large number such that the two fragments are non-interacting). The new datapoint keeps most of the properties of the original datapoints while respecting this change in coordinates. An additional “integer charge” loss term is added to the loss function to penalize far from integer charges on each fragment. Instead of training to the original DFT labels, we train these datapoints to MLIP labels (i.e. replacing dipole loss with partial charge loss) which are from a one shot forward pass while the augmented data is generated, using an existing model that has good test set performances (in this case, the MPNICE-RE model). The goal of the data augmentation is to ensure good distribution of the radical charge from single electron reduction, hence each augmented datapoint has a net charge such that one cluster is unreduced and the other (“infinitely separated”) cluster is unreduced.

MPNICE-RE-ISA-I only contains augmented datapoints where one cluster contains a bent EC and the other cluster does not, hence, the inference step to generate MLIP labels restricts the total charge to occur on the cluster with the bent EC. This calculation is effectively MLIP inference on the two separate fragments/clusters (respecting that the bent EC is reduced), and then combining these labels (energies, forces, charges) into an “infinitely separated” pair, ensuring training to size consistency.

Of the 20% augmented data in MPNICE-RE-ISA-II, 67% is the same procedure as MPNICE-RE-ISA-I. The remaining augmented data are pairs where both fragments contain a bent EC, and the MPNICE-RE, which was demonstrated to achieve an 0.73 kcal/mol RMSD on the vertical EA of bent clusters, is also used to determine which of these fragments has the more favorable vertical electron affinity. This pair still has the charge such that only one cluster is reduced, and the more favorably reduced cluster is MLIP-labelled in the reduced state. This method requires an initial model that predicts accurate vertical electron affinities, which the MPNICE-RE model does (see Table 1). Both ISA models were trained for 200 epochs each, transfer learning from the MPNICE-RE model which was trained for 400 epochs.

4.10 Visualization

Molecular structures are visualized with Schrödinger Materials Science Suite⁶⁷ and Vesta software.⁶⁸

5 Author Contributions

Y.W. performed DFT, MPNICE-MD, AIMD, dataset construction, MPNICE model training, validation/inference, ISA, analyzed data, developed software, and authored the manuscript. J.L.W. conceptualized the project, developed software, analyzed data, and extensively edited the manuscript. J.M.S. developed software and analyzed data. Z.K.G. provided concep-

tual input and edited the manuscript. X.X. developed software for ISA. L.D.J. conceptualized ISA. J.L.W., J.M.S., and R.A.F. supervised the project. All authors reviewed the manuscript.

6 Data Availability

Cluster benchmark sets will be made available upon publication at https://github.com/leifjacobson/MLFF_test_data.

Acknowledgement

We thank Garvit Agarwal and Rishabh Guha for valuable discussions and for QRNN-E MD data. The authors acknowledge support by Gates Ventures. The authors declare the following competing financial interest(s): R.A.F. has a significant financial stake in, is a consultant for, and is on the Scientific Advisory Board of Schrödinger, Inc.

Supporting Information Available

The following files are available free of charge.

- Supporting Information: Characterization of Off Center Radicals, Test Sets, Molecular Dynamics tests and analysis, DFT settings, Input File Examples (.pdf)

References

- (1) Behler, J. Perspective: Machine learning potentials for atomistic simulations. *The Journal of Chemical Physics* **2016**, *145*, 170901.
- (2) Deringer, V. L.; Caro, M. A.; Csányi, G. Machine Learning Interatomic Potentials as Emerging Tools for Materials Science. *Advanced Materials* **2019**, *31*, 1902765.

- (3) Butler, K. T.; Davies, D. W.; Cartwright, H.; Isayev, O.; Walsh, A. Machine learning for molecular and materials science. *Nature* **2018**, *559*, 547–555.
- (4) Jacobs, R. et al. A practical guide to machine learning interatomic potentials – Status and future. *Current Opinion in Solid State and Materials Science* **2025**, *35*, 101214.
- (5) Unke, O. T.; Chmiela, S.; Sauceda, H. E.; Gastegger, M.; Poltavsky, I.; Schütt, K. T.; Tkatchenko, A.; Müller, K.-R. Machine Learning Force Fields. *Chemical Reviews* **2021**, *121*, 10142–10186.
- (6) Wood, B. M.; Dzamba, M.; Fu, X.; Gao, M.; Shuaibi, M.; Barroso-Luque, L.; Abdelmaqsoud, K.; Gharakhanyan, V.; Kitchin, J. R.; Levine, D. S.; others UMA: A Family of Universal Models for Atoms. *arXiv preprint arXiv:2506.23971* **2025**,
- (7) Chen, B. W. J.; Zhang, X.; Zhang, J. Accelerating explicit solvent models of heterogeneous catalysts with machine learning interatomic potentials. *Chemical Science* **2023**, *14*, 8338–8354.
- (8) Schaaf, L. L.; Fako, E.; De, S.; Schäfer, A.; Csányi, G. Accurate energy barriers for catalytic reaction pathways: an automatic training protocol for machine learning force fields. *npj Computational Materials* **2023**, *9*, 1–10.
- (9) Schran, C.; Thiemann, F. L.; Rowe, P.; Müller, E. A.; Marsalek, O.; Michaelides, A. Machine learning potentials for complex aqueous systems made simple. *Proceedings of the National Academy of Sciences* **2021**, *118*, e2110077118.
- (10) Unke, O. T.; Stöhr, M.; Ganscha, S.; Unterthiner, T.; Maennel, H.; Kashubin, S.; Ahlin, D.; Gastegger, M.; Medrano Sandonas, L.; Berryman, J. T.; Tkatchenko, A.; Müller, K.-R. Biomolecular dynamics with machine-learned quantum-mechanical force fields trained on diverse chemical fragments. *Science Advances* **2024**, *10*, eadn4397.

- (11) Stevenson, J. M.; Jacobson, L. D.; Zhao, Y.; Wu, C.; Maple, J.; Leswing, K.; Harder, E.; Abel, R. Schrödinger-ANI: An Eight-Element Neural Network Interaction Potential with Greatly Expanded Coverage of Druglike Chemical Space. **2019**, arXiv:1912.05079 [physics].
- (12) Chen, M.; Jiang, X.; Zhang, L.; Chen, X.; Wen, Y.; Gu, Z.; Li, X.; Zheng, M. The emergence of machine learning force fields in drug design. *Medicinal Research Reviews* **2024**, *44*, 1147–1182.
- (13) Goodenough, J. B.; Park, K.-S. The Li-Ion Rechargeable Battery: A Perspective. *Journal of the American Chemical Society* **2013**, *135*, 1167–1176.
- (14) Xu, K. Nonaqueous Liquid Electrolytes for Lithium-Based Rechargeable Batteries. *Chemical Reviews* **2004**, *104*, 4303–4418.
- (15) Zhang, J.; Li, Y.; Chen, M.; Fu, J.; Zeng, L.; Tan, X.; Sun, T.; Feng, G. Atomistic simulation of batteries via machine learning force fields: from bulk to interface. *Journal of Energy Chemistry* **2025**, *106*, 911–929.
- (16) Dajnowicz, S.; Agarwal, G.; Stevenson, J. M.; Jacobson, L. D.; Ramezanghorbani, F.; Leswing, K.; Friesner, R. A.; Halls, M. D.; Abel, R. High-Dimensional Neural Network Potential for Liquid Electrolyte Simulations. *The Journal of Physical Chemistry B* **2022**, *126*, 6271–6280.
- (17) Wang, A.; Kadam, S.; Li, H.; Shi, S.; Qi, Y. Review on modeling of the anode solid electrolyte interphase (SEI) for lithium-ion batteries. *npj Computational Materials* **2018**, *4*, 1–26.
- (18) Ko, T. W.; Finkler, J. A.; Goedecker, S.; Behler, J. A fourth-generation high-dimensional neural network potential with accurate electrostatics including non-local charge transfer. *Nature Communications* **2021**, *12*, 398.

- (19) Ko, T. W.; Finkler, J. A.; Goedecker, S.; Behler, J. Accurate Fourth-Generation Machine Learning Potentials by Electrostatic Embedding. *Journal of Chemical Theory and Computation* **2023**, *19*, 3567–3579.
- (20) Jacobson, L. D.; Stevenson, J. M.; Ramezanghorbani, F.; Ghoreishi, D.; Leswing, K.; Harder, E. D.; Abel, R. Transferable neural network potential energy surfaces for closed-shell organic molecules: Extension to ions. *Journal of Chemical Theory and Computation* **2022**, *18*, 2354–2366.
- (21) Weber, J. L.; Guha, R. D.; Agarwal, G.; Wei, Y.; Fike, A. A.; Xie, X.; Stevenson, J.; Leswing, K.; Halls, M. D.; Abel, R.; Jacobson, L. D. Efficient Long-Range Machine Learning Force Fields for Liquid and Materials Properties. 2025; <https://arxiv.org/abs/2505.06462>.
- (22) Zubatyuk, R.; Smith, J. S.; Nebgen, B. T.; Tretiak, S.; Isayev, O. Teaching a neural network to attach and detach electrons from molecules. *Nature Communications* **2021**, *12*, 4870.
- (23) Ghasemi, S. A.; Hofstetter, A.; Saha, S.; Goedecker, S. Interatomic potentials for ionic systems with density functional accuracy based on charge densities obtained by a neural network. *Physical Review B* **2015**, *92*, 045131.
- (24) Zhang, L.; Wang, H.; Muniz, M. C.; Panagiotopoulos, A. Z.; Car, R.; E, W. A deep potential model with long-range electrostatic interactions. *The Journal of Chemical Physics* **2022**, *156*, 124107.
- (25) Kabylda, A.; Frank, J. T.; Dou, S. S.; Khabibrakhmanov, A.; Sandonas, L. M.; Unke, O. T.; Chmiela, S.; Müller, K.-R.; Tkatchenko, A. Molecular simulations with a pretrained neural network and universal pairwise force fields. **2025**,
- (26) Plé, T.; Adjoua, O.; Benali, A.; Posenitskiy, E.; Villot, C.; Lagardère, L.; Piquemal, J.-P. A foundation model for accurate atomistic simulations in drug design. **2025**,

- (27) Rappe, A. K.; Goddard, W. A. I. Charge equilibration for molecular dynamics simulations. *The Journal of Physical Chemistry* **1991**, *95*, 3358–3363.
- (28) Kocer, E.; Haouari, R. E.; Dellago, C.; Behler, J. Machine learning potentials for redox chemistry in solution. **2024**, arXiv:2410.03299 [physics].
- (29) Zubatiuk, R. AIMNet-NSE test datasets. 2021; <https://doi.org/10.5281/zenodo.5007980>.
- (30) Anusiewicz, I.; Skurski, P.; Simons, J. Fate of Dipole-Bound Anion States when Hydrated. *The Journal of Physical Chemistry A* **2020**, *124*, 2064–2076.
- (31) Barnes, T. A.; Kaminski, J. W.; Borodin, O.; Miller, T. F. I. Ab Initio Characterization of the Electrochemical Stability and Solvation Properties of Condensed-Phase Ethylene Carbonate and Dimethyl Carbonate Mixtures. *The Journal of Physical Chemistry C* **2015**, *119*, 3865–3880.
- (32) Timerghazin, Q. K.; Peslherbe, G. H. Non-nuclear attractor of electron density as a manifestation of the solvated electron. *The Journal of Chemical Physics* **2007**, *127*, 064108.
- (33) Timerghazin, Q. K.; Rizvi, I.; Peslherbe, G. H. Can a Dipole-Bound Electron Form a Pseudo-Atom? An Atoms-In-Molecules Study of the Hydrated Electron. *The Journal of Physical Chemistry A* **2011**, *115*, 13201–13209.
- (34) Young, R. M.; Neumark, D. M. Dynamics of Solvated Electrons in Clusters. *Chemical Reviews* **2012**, *112*, 5553–5577.
- (35) Bersuker, G. I.; Peng, C.; Boggs, J. E. The nature of the covalent bond: the existence and origin of nonnuclear attractors. *The Journal of Physical Chemistry* **1993**, *97*, 9323–9329.

- (36) Terrabuio, L. A.; Teodoro, T. Q.; Rachid, M. G.; Haiduke, R. L. A. Systematic Theoretical Study of Non-nuclear Electron Density Maxima in Some Diatomic Molecules. *The Journal of Physical Chemistry A* **2013**, *117*, 10489–10496.
- (37) Herbert, J. M.; Coons, M. P. The Hydrated Electron. *Annual Review of Physical Chemistry* **2017**, *68*, 447–472.
- (38) Stevenson, J. M.; Agarwal, G.; Bhai, L.; Debnath, S.; Goldsmith, Z. K.; Mondal, M.; Yeu, I.; Guha, R. D.; Svirinovsky-Arbeli, A.; Steingart, D.; others Evidence for significant multi-Li⁺ clustering in common lithium-ion battery electrolytes. **2025**,
- (39) Cazzanelli, E.; Croce, F.; Appetecchi, G. B.; Benevelli, F.; Mustarelli, P. Li⁺ solvation in ethylene carbonate–propylene carbonate concentrated solutions: A comprehensive model. *The Journal of Chemical Physics* **1997**, *107*, 5740–5747.
- (40) Bhatt, M. D.; Cho, M.; Cho, K. Interaction of Li⁺ ions with ethylene carbonate (EC): Density functional theory calculations. *Applied Surface Science* **2010**, *257*, 1463–1468.
- (41) Skarmoutsos, I.; Ponnuchamy, V.; Vetere, V.; Mossa, S. Li⁺ Solvation in Pure, Binary, and Ternary Mixtures of Organic Carbonate Electrolytes. *The Journal of Physical Chemistry C* **2015**, *119*, 4502–4515.
- (42) Müller, M.; Hansen, A.; Grimme, S. ω B97X-3c: A composite range-separated hybrid DFT method with a molecule-optimized polarized valence double- ζ basis set. *The Journal of chemical physics* **2023**, *158*.
- (43) Wang, Y.; Nakamura, S.; Ue, M.; Balbuena, P. B. Theoretical Studies To Understand Surface Chemistry on Carbon Anodes for Lithium-Ion Batteries: Reduction Mechanisms of Ethylene Carbonate. *Journal of the American Chemical Society* **2001**, *123*, 11708–11718.

- (44) Perdew, J. P.; Burke, K.; Ernzerhof, M. Generalized Gradient Approximation Made Simple. *Physical Review Letters* **1996**, *77*, 3865–3868.
- (45) Grimme, S.; Antony, J.; Ehrlich, S.; Krieg, H. A consistent and accurate ab initio parametrization of density functional dispersion correction (DFT-D) for the 94 elements H-Pu. *The Journal of Chemical Physics* **2010**, *132*, 154104.
- (46) Debnath, S.; Neufeld, V. A.; Jacobson, L. D.; Rudshiteyn, B.; Weber, J. L.; Berkelbach, T. C.; Friesner, R. A. Accurate Quantum Chemical Reaction Energies for Lithium-Mediated Electrolyte Decomposition and Evaluation of Density Functional Approximations. *The Journal of Physical Chemistry A* **2023**, *127*, 9178–9184, PMID: 37878768.
- (47) Agarwal, G.; Brock, C. N.; Guha, R. D.; Rao, K. K.; Stevenson, J. M.; Tiwari, S. C.; Fonari, A.; Jacobson, L. D.; Kwak, H. S.; Halls, M. D. Insights into Electrolyte Reactivity at the Li Metal Surface from Density Functional Theory. *ChemRxiv* **2025**,
- (48) Chen, J.; Martínez, T. J. QTPIE: Charge transfer with polarization current equalization. A fluctuating charge model with correct asymptotics. *Chemical Physics Letters* **2007**, *438*, 315–320.
- (49) Koski, J. P.; Moore, S. G.; Clay, R. C.; O’Hearn, K. A.; Aktulga, H. M.; Wilson, M. A.; Rackers, J. A.; Lane, J. M. D.; Modine, N. A. Water in an External Electric Field: Comparing Charge Distribution Methods Using ReaxFF Simulations. *Journal of Chemical Theory and Computation* **2022**, *18*, 580–594.
- (50) Ongari, D.; Boyd, P. G.; Kadioglu, O.; Mace, A. K.; Keskin, S.; Smit, B. Evaluating Charge Equilibration Methods To Generate Electrostatic Fields in Nanoporous Materials. *Journal of Chemical Theory and Computation* **2019**, *15*, 382–401.
- (51) Gergs, T.; Schmidt, F.; Mussenbrock, T.; Trieschmann, J. Generalized Method for

- Charge-Transfer Equilibration in Reactive Molecular Dynamics. *Journal of Chemical Theory and Computation* **2021**, *17*, 6691–6704.
- (52) Vondrák, M.; Reuter, K.; Margraf, J. T. Pushing Charge Equilibration Based Machine Learning Potentials to their Limits. **2025**,
- (53) Jensen, F. Unifying Charge-Flow Polarization Models. *Journal of Chemical Theory and Computation* **2023**, *19*, 4047–4073.
- (54) Krizhevsky, A.; Sutskever, I.; Hinton, G. E. ImageNet classification with deep convolutional neural networks. *Communications of the ACM* **2017**, *60*, 84–90.
- (55) Mazitov, A.; Bigi, F.; Kellner, M.; Pegolo, P.; Tisi, D.; Fraux, G.; Pozdnyakov, S.; Loche, P.; Ceriotti, M. PET-MAD, a universal interatomic potential for advanced materials modeling. *arXiv preprint arXiv:2503.14118* **2025**,
- (56) Stanton, R.; Kaymak, M. C.; Niklasson, A. M. N. Shadow Molecular Dynamics for a Charge-Potential Equilibration Model. *Journal of Chemical Theory and Computation* **2025**, *21*, 4779–4791.
- (57) Neese, F. The ORCA program system. *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **2012**, *2*, 73–78.
- (58) Chai, J.-D.; Head-Gordon, M. Systematic optimization of long-range corrected hybrid density functionals. *The Journal of Chemical Physics* **2008**, *128*, 084106.
- (59) Chai, J.-D.; Head-Gordon, M. Long-range corrected hybrid density functionals with damped atom–atom dispersion corrections. *Physical Chemistry Chemical Physics* **2008**, *10*, 6615–6620.
- (60) Rappoport, D.; Furche, F. Property-optimized Gaussian basis sets for molecular response calculations. *J. Chem. Phys.* **2010**, *133*, 134105.

- (61) Weigend, F.; Ahlrichs, R. Balanced basis sets of split valence, triple zeta valence and quadruple zeta valence quality for H to Rn: Design and assessment of accuracy. *Phys. Chem. Chem. Phys.* **2005**, *7*, 3297.
- (62) Ramezanghorbani, F. Developing Machine Learning Models to Enhance Applicability of Neural Network Potentials in Drug Discovery. Ph.D. thesis, University of Florida, 2020.
- (63) Cipolla, R.; Gal, Y.; Kendall, A. Multi-task Learning Using Uncertainty to Weigh Losses for Scene Geometry and Semantics. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2018; p 7482–7491.
- (64) Spotte-Smith, E. W. C.; Blau, S. M.; Xie, X.; Patel, H. D.; Wen, M.; Wood, B.; Dwaraknath, S.; Persson, K. A. Quantum chemical calculations of lithium-ion battery electrolyte and interphase species. *Scientific Data* **2021**, *8*, 203.
- (65) Giannozzi, P. et al. QUANTUM ESPRESSO: a modular and open-source software project for quantum simulations of materials. *Journal of Physics: Condensed Matter* **2009**, *21*, 395502.
- (66) Giannozzi, P. et al. Advanced capabilities for materials modelling with Quantum ESPRESSO. *Journal of Physics: Condensed Matter* **2017**, *29*, 465901, Publisher: IOP Publishing.
- (67) Schrödinger, LLC Maestro, Materials Science Suite. 2025; Schrödinger Release 2025-2.
- (68) Momma, K.; Izumi, F. VESTA 3 for three-dimensional visualization of crystal, volumetric and morphology data. *Journal of Applied Crystallography* **2011**, *44*, 1272–1276.

TOC Graphic

