White Paper: “Bridging the Gap: Transitioning from Deterministic to Stochastic Interaction Modeling in Electrochemistry”

IPAM Long Program, Fall 2025

Chris Anderson - UCLA
Ilgar Baghishov - University of Texas at Austin
Yagnik Bandyopadhyay - Arizona State University
Clotilde S. Cucinotta - Imperial College London
Sam Dillon - University of Florida
Sudheesh Kumar Ethirajan - University of California, Davis
Axel Groß - Ulm University
Christoph Freysoldt - Max Planck Institute for Sustainable Materials
Jürgen Fuhrmann - Weierstrass Institute, Berlin
Jonathan Gjerde - University of Michigan
Richard G. Hennig - University of Florida
Jirui Jin - University of Florida
Sung Hoon Jung - University of Texas at Austin
Jason Landini - University of Michigan
Lang Li - Fritz-Haber-Institut der Max-Planck-Gesellschaft
Mingjie Liu - University of Florida
Jörg Neugebauer - Max Planck Institute for Sustainable Materials
Vinh Nguyen- Michigan State University
Danny Perez- Los Alamos National Laboratory
Adam Petrucci - Michigan State University
Craig Plaisance - Louisiana State University
Keith Promislow - Michigan State University
Karsten Reuter - Fritz-Haber-Institut der Max-Planck-Gesellschaft
Stefan Ringe - Korea University
Soham Savarkar - University of Florida
Barbara Sumić - Fritz-Haber-Institut der Max-Planck-Gesellschaft
Katsuyo Thornton - University of Michigan
Mira Todorova - Max Planck Institute for Sustainable Materials
Ria Vasdev - University of South Carolina
Truong Vu - University of Illinois Chicago
Ping Yang - Los Alamos National Laboratory

Executive Summary

Electrochemistry bridges chemistry, materials science, biomolecular science, and mathematics to explain how electrical and chemical processes interact at the atomic scale. It underpins critical technologies—including batteries, fuel cells, electrolyzers, metal separation, and green-chemical synthesis—and provides a scientific foundation for sustainable technologies addressing renewable-energy integration, carbon reduction, and resource efficiency. Electrochemical interfaces also play a central role in biology, enabling molecular probing, diagnostics, and nanoscale bioelectronic technologies.

Despite this broad impact, electrochemical reactions occur in an environment that is exceptionally difficult to interrogate directly: the solid–liquid interface. Experiments typically access only macroscopic observables—applied voltages, ionic concentrations, electrolyte composition—while the microscopic structure and dynamics must be inferred indirectly. Deterministic continuum models, which describe averaged electrostatic potentials and ion distributions normal to an interface, historically provided the essential link between atomic-scale processes and signals measured by potentiostats, galvanostats, cyclic voltammetry, and impedance spectroscopy.

However, recent experimental and ab initio simulation advances reveal that electrochemical interfaces exhibit large statistical fluctuations in the electrostatic potential—in the order of one volt. These fluctuations profoundly influence molecular ordering, interfacial structure, adsorption, and reaction kinetics. They also challenge the assumptions underlying deterministic continuum descriptions, which smooth out precisely the spatial and temporal variations that govern interfacial chemistry. While ab initio simulations capture these fluctuations and provide mechanistic insight, they remain far too computationally intensive for the large spatial domains and high-throughput explorations needed to guide materials discovery and connect to experiment.

To move beyond these limitations, stochastically informed interaction-kernel models provide a scalable mathematical framework that retains the efficiency of continuum approaches while incorporating fluctuation statistics learned from atomistic simulations. By training interaction kernels on stochastically driven data, these models explicitly account for oscillatory water structure, fluctuating dielectric response, and rare but chemically decisive field fluctuations that deterministic models systematically miss. This stochastic, kernel-based modeling paradigm—successful in physics, biology, finance, and engineering—remains largely unexplored in electrochemistry but holds significant promise for unifying atomic-scale insight with macroscale predictability.

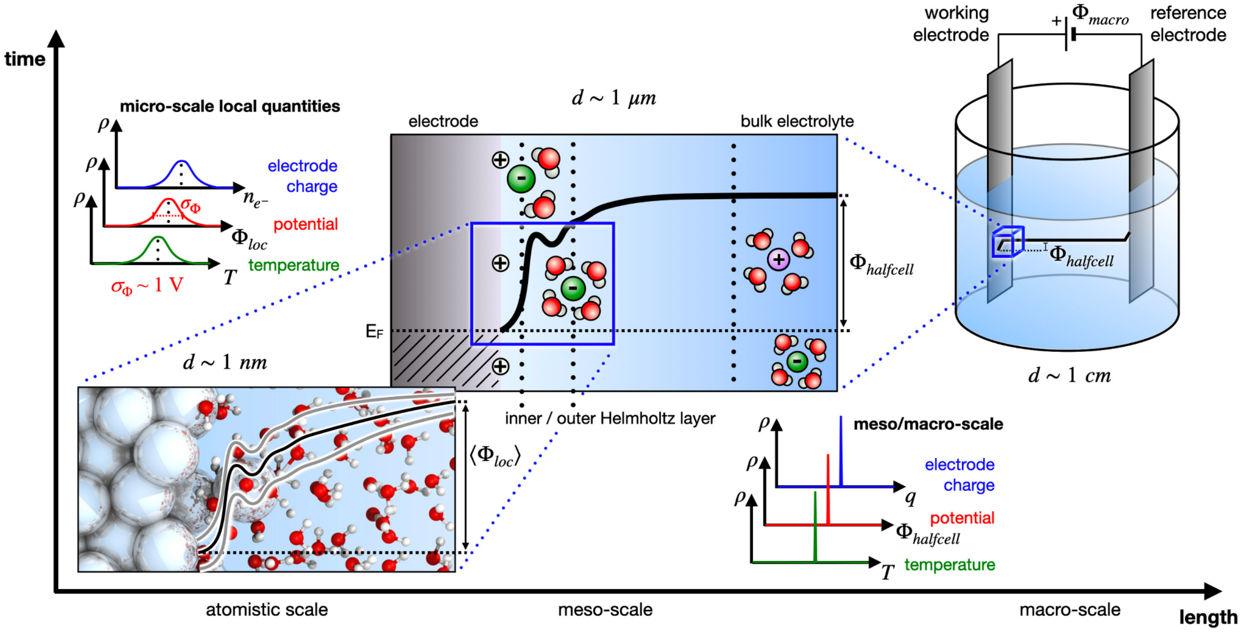

Fig.1: Schematic representation of an electrochemical system depicting (from left to right) the microscopic, mesoscopic, and macroscopic length- and time-scales, including the corresponding laterally averaged electrostatic potential perpendicular to the solid-liquid interface. Fluctuation-induced distributions of key electrochemical quantities like the electrode charge, electrostatic potential, and temperature are shown for the microscopic and meso-/macroscopic scales.

Section 1: Fluctuations in electrochemical systems

Background

Atomistic simulations of electrified interfaces consistently reveal strong spatial and temporal fluctuations of the electrostatic potential. For typical density functional theory (DFT) supercells at metal–electrolyte interfaces, electrostatic potential variations of the order of 1 V across the cell are observed.

In simple bulk systems with short-range interactions, statistical considerations based on the central limit theorem suggest fluctuations decay as $1/\sqrt{N}$ with system size. At electrochemical interfaces, this breaks down: the interface acts effectively as a layer of charge, and Poisson's equation implies that the associated field, and hence the potential, can remain correlated even over macroscopic distances. With periodic boundary conditions in-plane and open or asymmetric conditions along the surface normal, correlation lengths in water comparable to or larger than the lateral cell size mean that finite-size and image effects strongly shape the fluctuation spectrum of the electrostatic potential.

For electrochemistry, however, we are rarely interested in the electrostatic potential itself. What matters are free-energy differences along reaction coordinates (proton transfer, hydrogen evolution, adsorption, bond breaking) and their associated rates. Potential fluctuations enter through the coupling of local fields to charges and dipoles along these coordinates and via the entropy of the solvent and ions. Large apparent voltage fluctuations in a finite cell do not directly translate into eV-scale shifts in activation barriers; the energetic impact is controlled by $\Delta E \approx q\mathcal{E}\Delta z$, where Δz describes the displacement of the charge q along the reaction coordinate in the electric field $\mathcal{E}$.

Outstanding Challenges

Despite substantial progress, a number of conceptual and technical hurdles still limit our ability to treat potential fluctuations at electrochemical interfaces and to quantify their impact on various key quantities, such as effective barrier heights, in a quantitatively reliable way:

- **Quantifying fluctuation statistics.** Lack of information on the magnitude, spatial extent, and time scales of electrostatic potential and field fluctuations in bulk electrolyte and near realistic metal electrodes.

- **Ambiguity between physical vs. artificial fluctuations.** We must distinguish the impact of genuine thermal fluctuations dictated by the system's capacitance from spurious fluctuations or suppression of fluctuations induced by finite-size effects, boundary conditions, and specific algorithmic choices on computed properties.

- **Translating fluctuations to free energies and rates.** No general framework exists that allows for quantifying the impact of fluctuations on thermodynamics and kinetic parameters. There is a need to developing analysis methods that identify which observables depend only on free energies vs. those sensitive to the full fluctuation spectrum. Possible approaches are mapping potential fluctuations onto potentials of mean force (PMF) and rate constants.

- **Dynamic effects beyond static PMFs.** Memory, recrossing, and correlated "kicks" from the fluctuating medium may play an important role in the reaction rates, especially when slow modes of the environment are comparable to or slower than reaction times.

- **Finite-size scaling and asymptotic laws.** Asymptotic and scaling laws, including the dimensionality and anisotropy of electrostatic interactions at interfaces, are needed to make simulations at the thermodynamic limit tractable, but they are currently missing.

Work Done at IPAM

A core outcome of the IPAM program was the identification and clarification of a specific source of artificial potential fluctuations: the use of 2D-periodic boundary conditions in slab electrostatics. In standard slab geometries with 2D periodicity, the electrostatic solver admits an unphysical uniform in-plane $q_{||}=0$ mode that makes the plane-averaged potential along the perpendicular z direction execute a random walk; its variance diverges with slab thickness, in stark contrast to thermodynamic expectations. Analytical treatment of Poisson's equation and MD simulations of water show that this 2D-mode catastrophe enhances apparent fluctuations for small lateral cells, e.g., 1 nm^2. Control geometries with 1D and 3D periodicity do not show this divergence, confirming that the effect is a boundary-condition-induced artifact rather than an intrinsic property of the material. The MD simulations reproduce the insight from the thermopotentiostat concept, which links the amplitude of the intrinsic electrostatic potential fluctuations with the inverse of the capacitance of the system.

Within the broader IPAM electrochemistry program, we used these insights to organize thinking about how microscopic potential fluctuations should be interpreted and controlled. A first step was to establish a common language for “potential” in simulations: (i) global quantities such as work functions or cell-averaged potentials; (ii) local electrostatic potentials and fields near the reaction “hotspot”; and (iii) effective potentials projected along a chosen reaction coordinate. It is conceptually also of central importance to clearly distinguish between the locally varying electrostatic potential and the electrochemical potential that in thermodynamic equilibrium is constant throughout a cell.

Across different IPAM working groups, simulations consistently show electrostatic potential excursions of order ±1 V in typical DFT supercells, and the 2D-mode analysis and thermopotentiostat concept helped us distinguish which part of these excursions is genuine interfacial physics and which part is inflated by boundary conditions and finite-size artifacts.

From the thermodynamic side, participants converged on the PMF as the natural starting point for describing elementary electrochemical steps. PMFs obtained from constrained molecular dynamics (MD) or enhanced sampling provide thermodynamic barriers that incorporate the average effect of solvent and double-layer fluctuations.

Finally, the IPAM working groups formulated a roadmap:

- Develop computationally accurate methods to systematically quantify fluctuation magnitudes and correlation functions in realistic interfacial systems under potential control, and
- explicitly analyze and correct for 2D-mode-driven finite-size artifacts via alternative electrostatics, control geometries, and scaling laws.

This roadmap positions the community to move from informal awareness of potential fluctuations and boundary artifacts to a quantitatively controlled, method-spanning treatment of their effects on electrochemical thermodynamics and kinetics.

Scientific Opportunities

The clarification of how genuine interfacial fluctuations are entangled with electrostatic artefacts due to finite size effects opens up several concrete directions for progress, both conceptual and methodological. First, there is a clear opportunity to **design fluctuation-aware electrostatics for interfaces**. Guiding principles should be explored to develop and benchmark alternative kernels, boundary conditions, and post-processing corrections that suppress or analytically remove the offending mode, and to transfer these ideas systematically to other polar systems such as nanoporous electrodes, ionic liquids, and ferroelectric films.

Second, the program highlights the need for **quantitative fluctuation benchmarks and scaling laws** to study the impact of the potential fluctuations on physically relevant quantities such as reaction barriers. By combining analytical results with systematic numerical studies, one can derive practical criteria for lateral cell sizes, slab thicknesses, and screening lengths that ensure physically meaningful fluctuation spectra. These benchmarks can then be used to calibrate coarse-grained or implicit-solvent models that reproduce not just mean potentials and capacitances, but also the correct variance and correlation functions of the local electric fields that drive the reaction dynamics. Such models, in turn, can dramatically extend time and length scales, enabling statistically converged studies of rare events and slow collective modes that are currently out of reach for fully *ab initio* simulations.

Section 2: Potential control in stochastic systems

Background

Electrochemical experiments operate under **potential control**, where the electrode potential is fixed, and the resulting charge, structure, and reactivity are measured. In contrast, atomistic simulations use finite periodic cells whose effective capacitance is extremely small, leading to electrostatic potential fluctuations of order 1 V, as outlined in Sec. 1. This creates an inherent conflict for constant-potential algorithms. Fixing the potential in a small periodic cell suppresses natural fluctuations and distorts charge redistribution, double-layer structure, and reaction kinetics compared to an open-circuit macroscopic interface. Constant-charge simulations avoid this stiffness by allowing the potential to fluctuate, but then lose direct correspondence to the experimentally determined electrode potential and cannot model electrochemical, i.e. charge-changing reactions. Thus, nanoscale potential control requires not only setting the average electrode potential but also reproducing the correct fluctuation statistics of potential and charge. The dissipation-fluctuation theorem makes this explicit: the interfacial capacitance governs the variance of potential (or charge) fluctuations, just as heat capacity governs energy fluctuations.

A second complication is the separation of time scales. Electrons in a metal respond essentially instantaneously to electrostatic perturbations, whereas ions and solvent reorganize more slowly. Interfacial collective modes (e.g., ice-like water layers) may evolve even more slowly and with strong lateral correlations. As a result, interfacial potential fluctuations constitute a time-dependent, noisy environment that modulates barrier heights and recrossing dynamics—effects that cannot be captured by a static potential-of-mean-force picture.

Implementing potential control in a DFT supercell requires addressing three constraints:

1. **Adding and controlling excess electronic charge** at the electrochemical interface. The simulation must allow the electrode to acquire or release charge consistent with a specified potential.

2. **Providing a compensating countercharge** to keep the periodic supercell neutral. Neutrality can be maintained through implicit solvent models, background or sheet charges, structured counterelectrodes, or other embedding approaches.
3. **Allowing electrode and countercharge fluctuations dictated by the dissipation–fluctuation theorem.** To reproduce the correct temporal variance of potential and charge, the electrode charge—and its compensating countercharge—must be free to fluctuate. This requires electronically open boundary conditions that permit controlled exchange of electronic charge between the interface and an external reservoir.

Several thermodynamically consistent approaches for potential control now exist, including the Green's-function method of Otani, the computational Ne electrode, and recent Hairy-Probe/non-equilibrium Green's function (NEGF) DFT schemes.

Outstanding Challenges

While these methods provide a direct way to set up potentiostat-controlled electrochemical calculations in DFT supercells, several issues need to be addressed, specifically related to uncertainty quantification and how to link the results obtained in the atomistic supercells with experiments that take place at the meso-/macro scale.

- **Ensemble choice at finite size.** For realistic cell sizes, constant-charge, constant-potential, and thermopotentiostat ensembles are no longer equivalent; the impact of supercell size convergence on reaction energetics and kinetics must be quantified.
- **Electronic structure along reactions at fixed potential.** Thermodynamically consistent potentiostat approaches must be used to follow electronic structure, charge transfer, and bonding changes along a reaction path (e.g. thermopotentiostat, grand-canonical DFT/ constant-potential NEGF).
- **Transparent spatial and electronic boundaries.** Simulation cells and boundaries must be designed to mimic semi-infinite electrodes and electrolytes without spurious reflections of fields and fluctuations.
- **Voltage referencing.** Models must connect the internal potential of a finite simulation cell to experimental reference scales, recognizing that voltage is always a difference between the electrode and a reference.

Work Done at IPAM

Participants surveyed and evaluated a wide range of methods for charging supercells and modeling electrochemical interfaces, including Berry-phase polarization approaches; dipole corrections; jellium or sheet-charge strategies; explicit or continuum solvent models (`VASPsol`, `Environ`, ESM–RISM); constant-inner-potential DFT; grand-canonical DFT; Hairy-Probes DFT; fictitious-charge-particle methods; and the computational counter-electrode. These methods were analyzed in detail during the workshops and working-group discussions.

A major conceptual advance was the thermopotentiostat scheme, in which the electrode is coupled to a large external reservoir. The resulting Langevin-type dynamics reproduce both the mean voltage and the correct variance of potential and charge fluctuations—analogous to how a thermostat controls temperature and energy fluctuations. Electronic open-boundary approaches such as the computational Ne electrode or Hairy Probes provide computational microscopic realizations.

One working group focused on `VASPsol++`, Hairy-Probes DFT, and the computational counter-electrode methods. As part of the IPAM hackathon, a first implementation of the Hairy Probes open-boundary method was integrated into `VASPsol++` to enable electronic potential control directly within a widely used plane-wave DFT framework. This approach provides an alternative to the computational Ne electrode that has been implemented in `Vienna ab Initio Simulation Package` (`VASP`). At this stage, the implementation is a proof of concept: systematic testing, validation, and benchmarking against reference systems are still pending, and will be needed before releasing the code and preparing the associated paper.

The program also developed a clearer statistical-mechanical and electrostatic interpretation of potential fluctuations. Analysis of the $q_{||} = 0$ Fourier components of the electrostatic kernel revealed a direct connection between fluctuation amplitudes and the inverse vacuum capacitance. For typical atomistic cell areas, this yields voltage fluctuations on the order of ~1 V, consistent with numerical simulations. This clarified why conventional constant-potential algorithms behave too stiffly for small periodic cells, whereas constant-charge simulations preserve fluctuations but do not correspond to experimental control.

From this perspective, traditional constant-potential approaches for small cells were recognized as overly stiff: they clamp the potential and effectively remove the natural fluctuations, whereas constant-charge simulations preserve fluctuations but do not map cleanly onto experimental conditions and are unable to describe situations where excess charge is created, e.g. by electrochemical reactions.

The discussions also clarified a useful separation of roles between continuum and atomistic descriptions:

- **Mean potential profile:** Continuum electrostatics (Poisson-type models) describe the average potential drop, diffuse-layer charge distribution, and mapping to voltage scales.
- **Fluctuations around the mean:** Atomistic simulations with thermopotential control enforce the correct variance and dynamics of potential and charge. Open-boundary schemes supply the needed electronic reservoir.

Scientific Opportunities

The developments described above open several avenues for making potential-controlled simulations both more realistic and more predictive for electrochemistry. A first opportunity is the systematic quantification of finite-size and boundary-condition artifacts. Presently, there are no accepted simulation protocols that assess how supercell size, ensemble choice, or boundary treatment influence key observables such as electrochemical reaction barriers and rates. Obtaining these quantities requires extensive free-energy sampling over millions to billions of configurations and nanosecond trajectories, far beyond the accessible time and length scales of direct DFT. A new generation of machine-learning potentials that faithfully captures long-range electrostatics, interfacial capacitances, and realistic voltage fluctuations is therefore needed to generate benchmark datasets and establish size-converged reference results.

A second opportunity lies in multiscale integration. Continuum electrostatics can provide the mean potential drop across the interface, the charge distribution in the diffuse layer, and the mapping between applied voltage and interfacial charge. Atomistic simulations equipped with thermopotential control or electronically open boundaries can then resolve the detailed structure, correlations, and reaction pathways at the inner Helmholtz plane, including the correct statistics of potential fluctuations. Coupling these levels of description—passing continuum-derived mean potentials and capacitances to the atomistic region, and feeding back microscopic energetics and capacitances to refine the continuum model—offers a path toward voltage-dependent predictive models that can be compared directly with experimental observables.

Section 3: Capturing temporal fluctuations in continuum models

Background

Electrochemical phenomena are governed by dynamic interplay between the interfacial and bulk processes. In particular, there is a need for reduced models (continuum) that can capture the rich temporal structure of bulk water. Considerable computational effort is expended in relaxing molecular models of water to thermodynamic quasi-equilibrium as physically relevant initial conditions for MD simulations. Dynamic continuum models are distinct from equilibrium (time-independent) models for dielectric (zero-frequency) response for polar fluids. Dynamic continuum models are a possible model substitution for molecularly resolved bulk water.

The Bogoliubov–Born–Green–Kirkwood–Yvon (BBGKY) hierarchy is a dominant tool for the derivation of continuum models for plasma dynamics, such as the Boltzmann equation. There has been limited use of the BBGKY approaches in the electrochemical community. This is partially due to the intrinsic asphericity of molecules comprising polar fluids and their dynamic polarizability; parameterization of these features yield Vlasov systems expressed in high numbers of independent variables.

Outstanding Challenges

- **Faithful and tractable continuum models.** Models for the continuum setting must simultaneously capture structural organization of bulk water, incorporating effects of dipole fluctuations originating from molecular motion near thermodynamic equilibrium, and be computationally tractable.

- **Accurate and stable numerical methods.** Numerical challenges in the simulation of high-dimensional Vlasov-type equations must be resolved, admitting accurate solutions on timescales relevant to the electrochemical setting.

Work Done at IPAM

A canonical BBGKY hierarchy was derived from the molecular Stockmayer Hamiltonian. The Stockmayer Hamiltonian dynamics were regularized and rescaled to enable the application of the mean-field machinery. Crucially, appropriate rescaling preserves the ratio of molecular diameter to interparticle spacing and a finite total dipole moment. A mean field analysis established that the characteristic flow of the resulting Vlasov-Stockmayer system maps onto solutions of the Stockmayer-Hamiltonian flow. When particles are restricted to planar motion, the Vlasov-Stockmayer equation is a hyperbolic PDE in six independent variables. Taking hydrodynamic limits in linear velocity and assuming spatially-uniform density with zero macroscopic velocity reduces the system to four independent variables. Further restriction to a spatially constant external electric field and initial data uniform in one spatial direction allowed reduction to a probability density $f = f(y, \theta, \omega)$ in terms of the remaining spatial dimension, y, angular orientation, θ, and angular velocity, ω. Non-dimensionalization via comparison to classical water yields a time scale on the order of 50-100 femtoseconds, as is consistent with molecular simulations. The reduced Vlasov-Stockmayer model is expressed as

$$\partial_\tau f + \omega \partial_\theta f + (\gamma_1 G[f] - \gamma_2 E) \cdot n^{\perp}(\theta)\, \partial_\omega f = 0,$$

in terms of applied electric field *E*, and a nonlocal integral term $G[f]$ arising from long-range electrostatic interactions, the molecular director field $n = n(\theta)$, and two dimensionless parameters, γ_1 and γ_2, which balance electrostatic force against inertial force.

Two numerical schemes were designed and implemented for simulation of the reduced Vlasov-Stockmayer. The first approach utilized the finite-element-based package PRISMS-PF with an explicit streamline upwind Petrov-Galerkin stabilization to develop a prototype application. This was successfully implemented for the Vlasov-Stockmayer. The second approach exploits the rank-one vector-gradient tensor product structure of the system. The resulting procedure consists of a pseudo-spectral spatial discretization and exponential integration for time step advancement. The method has spectral accuracy and unconditional time stability and exhibits no numerical heating.

Scientific Opportunities

The model reduction points to several promising directions for future investigation. A systematic dimensional reduction of Vlasov systems is required that incorporates detailed molecular structure. The development of concomitant computational techniques is needed to resolve the Vlasov-Stockmayer system in its original high dimensional form. A key step is the identification of time scale dichotomies which can be exploited to obtain reduced models. The accuracy of the Vlasov-Stockmeyer system can be improved by identifying higher-correlation corrections obtained through machine learning trained simulations. The training data would be optimized by selecting initial data for the Stockmayer-Hamiltonian flow sampled from N-particle data that is far from the chaotic regime.

The mean-field models must be validated against electrochemical data from molecular simulations. This includes the comparison of relevant time scales of ion transport to those of the major structural reorganization of bulk water and capturing frequency dependence of dielectric response.

The numerical schemes can be improved by boasting the temporal accuracy via higher-order Trotter product formulas. The impact of electrochemical interfaces can be examined via use of non-periodic boundary conditions implemented and singular electrostatic potentials using the finite element method.

Section 4: Coupling atomistic and continuum models

Background

Continuum models are a powerful tool for the study of electrochemical systems. Due to the reduced number of parameters they require, continuum models are more computationally tractable than molecular and ab initio models, allowing them to be scaled to device dimensions. Continuum models are readily transferable across electrolyte and material-composition space, yielding data that can be directly compared to experimental measurements. However, it is often difficult to justify the physical basis of these models, as differing assumptions and parameter choices can yield similar macroscopic predictions. This ambiguity has made it challenging to evaluate the extrapolative capacity of continuum models beyond canonical applications. Additional information, including experimental measurements and atomistic simulation data, is required to evaluate, refine, and ultimately enhance the transferability of continuum models.

A key goal is to resolve the structure and properties of the electric double layer (EDL). This structure is influenced by atomistic details, which can be challenging to observe experimentally. Ab initio atomistic simulations of the EDL can provide precisely the kind of detailed information required to understand the detailed physical phenomena that occur in the EDL. Ab initio chemical methods are typically unable to reach the long timescales required for EDL relaxation or the large spatial extent of the diffuse layer. Atomistic or machine-learned force fields can access the necessary timescales, but they introduce many parameters which can limit their transferability. Despite these challenges, such simulations yield valuable insights into the electrochemical interface on a case-by-case basis, offering benchmarks for developing and validating new continuum models.

Outstanding Challenges

- **Diverse physically relevant benchmarks.** There is a lack of well-defined benchmark problems that compare molecular dynamic simulation results and outcomes of various continuum models for a diverse collection of electrolytes.

- **Validation and improvement of continuum models for EDL.** Improved models should enable the quantitative reproduction of experimentally observed phenomena like limited double layer capacity and at the same time incorporate effects seen in molecular simulations like charge oscillations and layering of ions.
- **Bridging length and timescales**. Continuum models reach device-level dimensions, while atomistic simulations capture molecular fidelity; however, no existing framework seamlessly connects the two.
- **Incomplete sampling of EDL dynamics**. Atomistic simulations struggle to reach the long relaxation times or large diffuse-layer thicknesses needed for full equilibration, restricting their ability to reproduce experimentally observable electrostatic screening and ion distribution.

Work Done at IPAM

The goal was to develop strategies for bridging continuum and atomistic simulation approaches. A central challenge is the consistent matching of the simulation cell and electrochemical conditions, which are often treated disparately in the two approaches. By developing and implementing continuum and atomistic models of the electrochemical interface in close synergy—and by carefully aligning these variables—we compared their microscopic predictions, identified potential limitations of continuum models, and outlined pathways for improving them.

On the atomistic side, we constructed EDL model systems using `pyiron` workflows and performed classical molecular dynamics simulations under controlled surface charges. These simulations provided structural and electrostatic descriptors, including water, ion distributions, and potential profiles; however, limited by system sizes and timescales, only the early stages of interfacial equilibration were captured. Continuum modeling was based on a statistical mechanics approach to derive a free energy model for charged mixtures and nonequilibrium thermodynamics to derive phenomenological partial differential equations. The resulting generalized Nernst-Planck-Poisson system describes finite ion size and solvation effects, and includes the momentum balance, which equilibrates the pressure gradient with the Coulomb force. An ion conservation constraint was added to the generalized Poisson-Boltzmann (PB) problem describing the diffusional equilibrium of this system in order to reflect the structure of the closed symmetric cell. In addition, the dielectric decrement was described by a concentration and electric field dependent polarization free energy modeling the interaction of dipoles in an electric field, parametrized by experimentally known bulk susceptibility data. To be consistent, contributions to chemical potentials and to the momentum balance were taken into account. The resulting augmented PB problem was implemented in a Julia package `AugmentedPoissonBoltzmann.jl` using a finite volume approach and taking advantage of automatic differentiation for handling the highly nonlinear system. A python wrapper was developed to deploy this code within a `pyiron` workflow. This allowed reproducible and straightforward workflows and easy comparison between continuum and atomistic results.

Previous research of IPAM participants was extended to develop a model and derive analytical solutions for the charge distribution across an electrode grain boundary. The fundamental assumption of the model is that the differing work functions of constituent materials give rise to a potential difference across grains. The analytical solution showed that there is an accumulation of charge at the grain boundaries. When placed in an electrolyte, this inhomogeneous surface charge distribution may enhance the corrosion rate at the grain boundaries, explaining experimental observations.

Scientific Opportunities

Integrating atomistic simulations with continuum models offers a physically rigorous approach that surpasses traditional modeling techniques. The unification of the two approaches will foster an improved understanding of phenomena from interfacial chemistry and material behavior in experiments and technologies. The enhanced predictive accuracy will accelerate the development of corrosion-resistant metals and advance electrochemical energy storage systems.

Immediate priorities include developing robust coupling strategies between atomistic and continuum simulations—an initiative out of IPAM—and building open-source software platforms supporting collaborative workflows, such as `pyiron`. Examples of this include finite volume methods in Julia and finite element methods in `PRISMS-PF`, that support code implementation and verification. Benchmarking against experimental datasets will be vital for validating results and ensuring reproducibility. Modeling advances are needed for solvents beyond water and for ionic liquids. Incorporating these advances necessitate extensible software infrastructure across diverse systems.

While tightly coupled atomistic and continuum models offer high accuracy, their computational cost is substantial. Yet, once validated, these frameworks can generate valuable datasets for machine learning, enabling efficient materials screening via physics-aware surrogate models. For example, the incorporation of microstructural features such as grain orientations and grain boundaries will yield models with the ability to accurately describe observed material phenomena and enable identification of effective materials for electrochemical cells.

Section 5: Implicit electrolyte models for quantum simulations

Background

Accurate modeling of solid-liquid interfaces is essential for energy technologies. However, explicit solvent ab initio molecular dynamics simulations are still too computationally intensive. Implicit continuum models are an efficient alternative, but they often fail at charged interfaces. Conventional implementations neglect non-local responses and Pauli repulsion, resulting in a "dielectric catastrophe" in which screening occurs unphysically close to the surface. This distorts capacitance and solvation predictions, especially in regions dominated by Stern layer physics. We address these issues by decoupling the cavity geometry from the electronic response and integrating non-local physics. This allows us to capture realistic electrolyte gaps without the cost of explicit sampling.

Outstanding Challenges:

- **Decoupling error contributions.** Isolating the error sources in solvation free energies derived from geometric cavity definitions versus those arising from electronic dielectric response.
- **Interfacial transferability.** Addressing the rigidity of current cavity and response definitions to accommodate diverse solvent classes and the unique physics of the electrode-electrolyte interface.
- **Systematic error correction.** Identifying implementation-specific systematic errors across solvent classes to provide physics-based correction factors for heterogeneous interfaces, which differ fundamentally from bulk environments.

Work Done at IPAM

We integrated Effective constitutive elements from established models — specifically `SaLSA` (`JDFTx`), `VASPsol++`, and `pyiron` — to address Pauli repulsion and over-screening at electric interfaces. The primary objective was to decouple the geometric cavity definition from the electronic response. We determined accurate cavity creation to be more critical than non-linear dielectric response for correcting the unphysical screening onset at electrode interfaces.

We developed a Python prototype using a convolution-based shape function and non-local dielectric response. While non-self-consistent results were promising, the external implementation caused Hamiltonian discontinuities due to mismatches with `VASP`'s Projector Augmented Wave (PAW) one-center terms. Consequently, we migrated the method directly into the `VASP` Fortran source code to ensure consistent potential projection into PAW spheres and robust self-consistent field convergence.

We validated the implementation by reproducing aqueous solvation energies for water and the Minnesota Solvation Database, and by benchmarking against `JDFTx`. This analysis successfully decoupled error contributions arising from geometric cavity definitions versus polarization response.

To enable realistic modeling of electrochemical interfaces on standard periodic grids, we implemented two critical features. First, to capture semi-infinite diffuse layers without prohibitive computational cost, we applied a non-uniform coordinate mapping that approximates a logarithmic distance scale. This allowed the simulation to match the exponential decay of the diffuse layer while remaining on `VASP`'s standard spectral grid. Secondly, standard vacuum dipole corrections fail in implicit solvents because the medium responds to the correction field. To address this, we built a smooth dipole sheet directly into the non-linear Poisson–Boltzmann solver, ensuring that the electrolyte response and dipole correction converge self-consistently to eliminate spurious periodic interactions.

Scientific Opportunities

Integrating `SaLSA` and non-local response kernels into `VASP` enables high-fidelity implicit solvent modeling of solid-liquid interfaces by resolving the dielectric catastrophe and accurately capturing Pauli repulsion. This framework allows for precise simulation of the electrochemical double layer, essential for studying voltage-dependent phenomena like solid-electrolyte interface formation and electrocatalytic mechanisms. By successfully decoupling geometric cavity errors from electronic response limitations, this work opens a pathway for systematic model improvements and the derivation of transferable correction factors. This capability facilitates high-throughput screening of electrolytes with an accuracy previously restricted to computationally expensive explicit methods, ultimately democratizing access to predictive ab initio design of interfacial materials.

A key opportunity identified during the program is the development of an Implicit Quantum Electrolyte: a framework that introduces quantum-mechanical coupling between the explicit system and the implicit solvent. This approach should retain $O(N \log N)$ scaling through convolution-based operators and fast Fourier transforms, enabling quantum-consistent solvent effects — nonlocal structure, directional alignment, hydrogen bonding, and charge transfer — without explicit MD.

Section 6: Generative Artificial Intelligence

Background

Generative artificial intelligence (AI) has become a powerful tool for accelerating electrochemical discovery, providing a data-driven alternative to costly simulation methods such as DFT, MD, and materials screening. By learning structural and energetic distributions directly from data, these models can rapidly propose new electrode materials, generate physically realistic solvent configurations around the EDL, or model voltage-dependent structural states. Recent advances in autoregressive, diffusion, and flow-matching frameworks demonstrate the ability to capture complex free-energy landscapes and key electrochemical observables such as e.g. local potential, charge densities, significantly reducing computational effort. However, important challenges remain, including handling multi-ionic environments, ensuring energetic and geometric stability, and incorporating physically meaningful guidance to maintain chemical validity.

Outstanding Challenges

- **Energetic stability.** Ensuring energetic stability of generated electrodes or liquid configurations while minimizing the need for extensive DFT or MD simulations.
- **Voltage decorrelation.** Mitigating Li vs. Na voltage correlation artifacts that limit controllable property-guided generation.
- **Electrochemical constraint.** Incorporating constraints specific to electrochemistry that go beyond the geometric invariance or equivariance of standard graph neural networks

architectures, e.g. interfacial charge distribution, solvation shell structure, electrochemical potential compatibility.

- **Target-conditioned generation.** Building conditional generative pipelines that respond accurately to imposed targets, e.g. Li-voltage in the water stability window, Li and Na voltage difference.
- **Benchmarking.** Developing unified benchmarks of different generative models to evaluate chemical validity, stability, and functional performance of generated structures.

Work Done at IPAM

In electrode and electrolyte discovery, generative models must balance chemical validity, thermodynamic stability, and functional properties such as voltage and selectivity. Models like `MatterGen`, guided diffusion, `Universal Model for Atoms` (`UMA`), and `OMatG` have been used to generate Li-selective electrode structures by leveraging large cathode datasets and ion-swapping or property-guided sampling. While the generated materials show strong novelty, voltage correlations between Li and Na limit the ability to target high Li voltage and low Na voltage simultaneously, motivating strategies such as training on voltage differences rather than absolute values. A similar challenge appears in generative modeling of solvents like water, where capturing hydrogen-bond networks and realistic bond geometries is essential. Diffusion and autoregressive models can reproduce radial distribution functions (RDF) well. Yet, small structural artifacts such as distorted angles that can lead to unstable energies after relaxation. Across both problems, a common theme emerges: generative models learn structural distributions effectively, but without physics-based constraints or energetic guidance, they may generate plausible-looking yet chemically invalid configurations.

During the IPAM program, one subgroup explored two very different but complementary applications of generative AI within electrochemistry.

(1) Generative Design of Li-Selective Electrodes.

Using datasets of about 2,400 Li-battery and about 300 Na-battery cathodes, we fine-tuned several generative frameworks including `MatterGen`, `GuidedMatDiffusion`, `OMatG`, and `UMA`-based voltage predictors to propose new electrode materials suitable for Li-selective electrochemical separation. We evaluated assumptions such as swapping Li/Na/Mg ions and examined the reliability of generative voltage guidance. Although direct guidance was limited by correlations between Li and Na voltages in the dataset, we demonstrated that generative sampling combined with post-processing (local relaxation, ion-substitution heuristics) can identify stable, novel candidates with desirable voltage profiles.

(2) Generative Modeling for Accelerated Water Sampling.

In a separate track, we investigated generative approaches for fast sampling of bulk water configurations which is an essential component for electrochemical interface modeling. We trained both an autoregressive model and a diffusion model to reproduce water RDFs and structural distributions based on MD snapshots. Autoregressive sampling was exponentially slow with increasing system size, while diffusion produced bulk water structures with 1,000 molecules in minutes. These models highlight the potential for generative AI to reduce MD sampling requirements dramatically when paired with physics-consistent constraints.

Together, these two efforts of materials generation and molecular-scale sampling illustrate the breadth of generative AI's impact across electrochemical science, from designing new solid electrodes to characterizing liquid structures at interfaces.

Scientific Opportunities

The complementary nature of the two projects points toward several high-impact scientific opportunities where generative AI could transform electrochemistry. Physics-guided (e.g., energy-guided) models by integrating force-fields, machine learning interatomic potentials (MLIP), or *ab initio* molecular dynamics (AIMD)—offer a pathway toward generating chemically valid structures with controllable voltage or stability targets. A unified generative framework that couples electrode-design models with water-sampling diffusion models could enable fully generative simulations of solid–liquid interfaces, accelerating the discovery of ion-transport mechanisms and reaction pathways. Conditional generative models, informed by properties such as ion radius, charge state, coordination environment, or redox potential, may systematically produce materials optimized for selectivity and reactivity in separations, batteries, or catalysis. Moreover, generative models can accelerate MD and DFT workflows by supplying high-quality initial configurations or providing rich training data for MLIPs, reducing computational cost while accessing larger time and length scales. Developing cross-domain latent representations that integrate structural, energetic, and dynamical information could further unlock exploration of chemical spaces unreachable through conventional simulations. Finally, generative-AI–driven discovery of electrolyzers presents a promising direction, enabling the design of systems with expanded electrochemical windows and tailored properties that surpass current aqueous technologies limited by a narrow ~1.2 eV redox range.

Section 7: Electrostatics in machine learned interatomic potentials

Background

MLIPs are a recent breakthrough in atomistic modeling. They are essential for bridging the gap between the highly accurate but computationally costly DFT calculations that explicitly account for all electrostatic interactions and empirical force-field modeling, where the electronic degrees of freedom are coarse-grained away as they enable simulating large system sizes on a nanosecond timescale with near-DFT accuracy.

However, they rely on the assumption that the total energy of a system can be decomposed to atomic contributions on a short-range order within a defined cutoff radius. The assumption of locality omits some important interactions that are long-ranged and electrostatic in nature. The dipole-field interaction, for example, governs the structuring of the solvent under an applied electric field, either due to the electrode or due to ions present in the electrolyte. The Coulombic interaction between ions, as well as between ion and electrode (image charge effect) affects the local fields and the solvent organization geometry, and is long-ranged. Furthermore, charge transfer at the interface drives all electrochemical reactions. Thus, it is essential to capture charge transfer in order to study electrocatalytic processes.

The common approach of electrostatics-aware MLIPs is to include an explicit Coulomb-law expression to capture long-range electrostatic effects. The various approaches differ, however, in how the charges (and possibly dipoles/higher multipoles) are determined. This architectural design translates into which of the specific phenomena listed above are captured.

Outstanding Challenges

- **Applicability range of ML electrostatic models.** Clear guidance for practitioners on the suitability of a particular machine-learning electrostatic model for a specific application is presently lacking. Establishing a standardized benchmark, accompanied by a public leaderboard on which each newly proposed electrostatics aware ML potential is evaluated, would enable informed model selection.
- **Standardized electrostatics benchmarks.** Electrostatic phenomena are diverse (for examples, see below) and thus impose distinct requirements on the models attempting to reproduce them. At present, the community lacks a standardized suite of electrostatic benchmark problems that could be used to assess the applicability of a given electrostatics-aware MLIP to each of them.
- **Availability and interoperability.** In one IPAM working group, we observed that, from an application perspective, code accessibility and usability are major barriers. MLIP codes, especially those with electrostatic functionality, should be distributed as ready-to-use packages that guarantee seamless interoperability with established atomistic simulation suites, such as LAMMPS and GROMACS, and with reproducible environments, such as MyBinder. Direct access to internal charges and coupling to external fields/potentials are necessary for electrochemical applications. This strategy will promote broader adoption, reproducibility, and integration into existing computational workflows.

Work Done at IPAM

The working group focused on assessing and comparing how current electrostatics-aware MLIP schemes capture the long-range and non-local electrostatic effects that are fundamental to electrochemical systems. We identified a set of representative physical phenomena that any electrostatics-aware MLIP must be capable of reproducing. These include the image effect arising from charge redistribution near metallic or dielectric interfaces, the inhomogeneous distribution of excess charge on rough conductive surfaces, and position-dependent charge transfer associated with Fermi-level pinning. The correct treatment of environment-dependent molecular dipoles, electronic polarization, and the magnitude and dynamics of polarization fluctuations in water – and how they are affected by local fields – is also essential, as these collectively determine screening, field-induced forces, and the thermodynamic response of polar liquids.

After surveying the literature and examining the MLIP architectures, we identified four approaches through which long-range electrostatics are incorporated into otherwise short-range models. Latent charge schemes are self-supervised models that learn internal environment-dependent charges from a separate neural network alongside the short-range potential. Wannier-centroid schemes train an auxiliary model on dipoles or Wannier centers and compute long-range electrostatics analytically on top of the short-range model. Charge-equilibration (QEq) schemes obtain atomic charges by minimizing an electrostatic functional under a fixed-charge constraint, with environment-dependent electronegativities and hardness parameters. Response-learning approaches directly supervise the model on response properties, such as field derivatives of the energy or polarizabilities, and incorporate these into the total energy expression.

To evaluate the ability of existing models to capture long-range electrostatics, we constructed a test system consisting of polar water slabs. We then computed the potential energy profile obtained by translating either a hydronium or hydroxide ion along the slab normal.

To determine which long-range contributions the ML model correctly reflects and which it lacks, we implemented a continuum-level reference framework in pyiron with adjustable parameters that control dielectric response, slab polarization, and background charge compensation. Our analysis revealed that, in the absence of an explicit charge-neutrality constraint, some models redistributed charge across the neutral water slab, effectively introducing a compensating background charge that altered the energetic response. Other models more accurately captured dipole formation within the slab, producing energy profiles characteristic of a dielectric or capacitor-like system.

This comparative analysis enabled us to identify the shortcomings of existing charge-aware ML models, which is a critical step toward developing the next generation of models that are free of these artifacts.

Scientific Opportunities

An outstanding challenge is the lack of a well-defined test suite for prototypical phenomena that electrostatics-aware MLIP should be able to reproduce. The test-suite developed in our working group provides a comprehensive set of fundamental mechanisms reflecting each specific qualitative aspect of long-range electrostatic interactions. Not all phenomena are equally relevant for all use cases, and the purpose is to guide developers and practitioners into what a specific model captures. A possible extension, more geared towards the electrochemical community, would be benchmarks for the complex systems encountered there (realistic electrode + adsorbate + water + ion environments), together with explicit system target properties such as capacitances, potentials of zero charge.

A community-maintained leaderboard and evaluation framework can host these benchmark datasets, supply fixed training/validation splits, and automatically compute metrics of physical consistency, predictive accuracy, and computational speed, enabling transparent comparison of electrostatics-aware MLIPs.