POLITECNICO DI TORINO
Repository ISTITUZIONALE

Electrified metal–water interfaces at constant potential: Data-efficient transfer learning for machine-learning-based molecular dynamics

Original
Electrified metal–water interfaces at constant potential: Data-efficient transfer learning for machine-learning-based molecular dynamics / Bianchi, M.G., Re Fiorentin, M., Risplendi, F., Pirri, C., Parrinello, M., Bonati, L., Cicero, G.. - In: ESCIENCE. - ISSN 2667-1417. - ELETTRONICO. - 7:1(2027), pp. 1-10. [10.1016/j.esci.2026.100642]

Availability:
This version is available at: 11583/3015703 since: 2026-10-01T12:15:25Z

Publisher:
Elsevier

Published
DOI:10.1016/j.esci.2026.100642

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

28 November 2026

eScience 7 (2027) 100642

Contents lists available at ScienceDirect

eScience

journal homepage: www.keaipublishing.com/en/journals/escience

Research Paper

Electrified metal–water interfaces at constant potential: Data-efficient transfer learning for machine-learning-based molecular dynamics

Michele Giovanni Bianchi[a,b], Michele Re Fiorentin[a], Francesca Risplendi[a], Candido Fabrizio Pirri[a,c], Michele Parrinello[b], Luigi Bonati[b,*], Giancarlo Cicero[a,*]

[a] Department of Applied Science and Technology, Politecnico di Torino, Corso Duca Degli Abruzzi 24, Torino, 10129, Italy
[b] Atomistic Simulations, Italian Institute of Technology, Via Enrico Melen 83, Genova, 16152, Italy
[c] Centre for Sustainable Future Technologies, Italian Institute of Technology, Via Livorno 60, Torino, 10144, Italy

HIGHLIGHTS

- A workflow for ML interatomic potentials at electrochemical interfaces.
- Transfer and active learning strategies boost data efficiency in ML models.
- Efficiency enables high-level ground-truth and constant-potential schemes.
- ML models for copper/water interfaces are trained on only a thousand data points.
- ML reveals bias-dependent solvent restructuring phenomena at electrified surfaces.

GRAPHICAL ABSTRACT

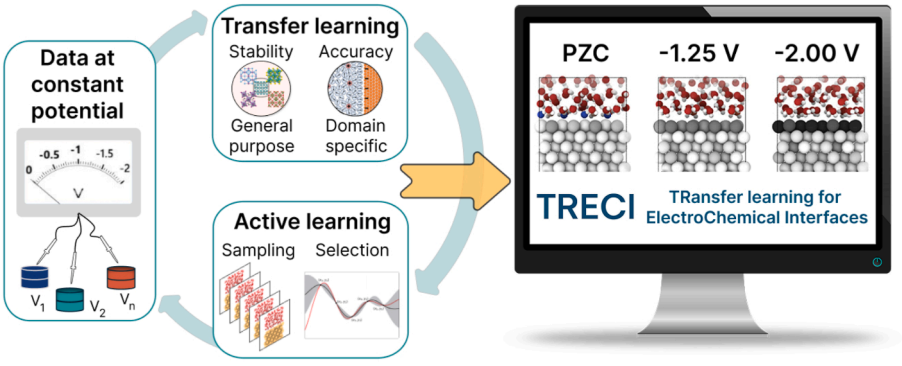

ARTICLE INFO

Keywords:
Electrified metal–water interfaces
Constant potential
Machine learning force-field
Transfer learning

ABSTRACT

Simulating electrified metal–water interfaces with explicit solvent under constant potential is essential for understanding electrochemical processes, yet remains prohibitively expensive with *ab initio* methods. We present TRECI, a data-efficient workflow for constructing machine learning force-fields (ML-FFs) that achieve *ab initio*-level accuracy in electronically grand-canonical molecular dynamics. By leveraging transfer learning from general-purpose and domain-specific models, TRECI enables stable and accurate simulations across a wide potential range using a reduced number of reference configurations. This efficiency allows the use of high-level meta-GGA functionals and rigorous surface-electrification schemes. Applied to Cu(111)–water, models trained on just one thousand configurations yield accurate molecular dynamics simulations, capturing bias-dependent solvent restructuring effects not previously reported. TRECI offers a general strategy for characterising diverse materials and interfacial chemistries, significantly lowering the cost of realistic constant-potential simulations and expanding access to quantitative electrochemical modelling.

*Corresponding authors.
E-mail addresses: luigi.bonati@iit.it (L. Bonati), giancarlo.cicero@polito.it (G. Cicero).
Peer review under the responsibility of Editorial Board of eScience.

https://doi.org/10.1016/j.esci.2026.100642
Received 24 December 2025; Received in revised form 18 May 2026; Accepted 1 September 2026
Available online 5 September 2026
2667-1417/© 2026 The Authors. Publishing services by Elsevier B.V. on behalf of KeAi Communications Co. Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

1. Introduction

Charged metal–H_2O interfaces play an important role in fields such as corrosion, heterogeneous catalysis, and electrochemistry [1,2]. Despite decades of study, recently intensified by the rise of electrochemical technologies for the energy transition, understanding these interfaces at the atomic level remains a major challenge in surface science [3–9]. *Ab initio* molecular dynamics (AIMD) simulations are an ideal candidate to provide atomistic insight in these systems [10]. However, the heterogeneity of the interfaces composed of metal surfaces, solvents and solutes clashes with the problem of small spatial and short time scales achievable in AIMD. In recent years, machine learning interatomic potentials, here referred to as ML force-fields (ML-FFs) to avoid confusion with the applied electric potential, have emerged as a promising solution to overcome the limits of AIMD. They bridge the accuracy of *ab initio* approaches with the computational feasibility of force-field schemes [11]. However, the presence of an applied electrical bias significantly increases the simulation complexity both at the *ab initio* and ML level. The simplest approaches for the electrification of the interface are based on constant charge schemes [12], which are not fully representative of the experiments, since the applied bias is not controllable. In principle, constant potential methods align more closely with experimental ensembles [13]. However, they are more computationally expensive, and in practice they are applied in conjunction with implicit solvent schemes [14–16] which can introduce artefacts.

In addition to these issues associated with the electrification schemes, ML-FFs face a further complication: the treatment of long-range electrostatic interactions. The presence of charged species can give rise to interactions that extend beyond the local nature of the force-fields [17,18]. As a result, the development of ML-FFs for electrochemical systems constitutes an area of research that remains largely unexplored. Preliminary ML-FFs in the electrochemistry domain neglect the effects of the applied bias, focusing only on the potential of zero charge (PZC), i.e., when there is no extra charge on the surface [19–21]. To overcome these limitations, ML-FFs employing a constant charge approach were developed [9], while subsequent advances proposed constant potential ML models employing hybrid explicit-implicit solvent schemes [22–25].

In this work, we present TRECI, a data-efficient strategy that uniquely leverages transfer learning [26,27] and active learning to construct ML-FFs tailored for electrified metal–water interfaces. By featuring a fully explicit solvent model and a rigorous constant potential scheme (the "double reference method" [28]), TRECI allows for higher fidelity simulations, with a level of accuracy that was not possible before. We apply TRECI to a copper–water interface, a representative electrode used in many electrocatalytic applications [7,29]. TRECI achieves stable and accurate simulations with only a thousand configurations, a significant reduction of the dataset size in comparison to similar models trained from scratch. Notably, TRECI enables direct observation of bias-dependent solvent restructuring phenomena, offering insights that were inaccessible with prior models. The TRECI code is openly available at https://github.com/michelegiovannibianchi/TRECI, and its versatility allows straightforward application to other electrochemical systems. This paves the way for broader access to advanced constant potential simulations in a wide range of electrochemical applications.

2. Results and discussion

2.1. Data-efficient workflow for electrified metal–water interfaces

The main outcome is the development of TRECI (TRansfer learning for ElectroChemical Interfaces), a computational workflow for ML-FFs that provides a high-fidelity description of electrified metal-water interfaces while remaining data-efficient. The success of TRECI relies on three core pillars: 1) a high-quality quantum mechanical description of the interface, 2) a data-efficient transfer learning framework for building the ML-FFs and 3) an effective active learning strategy for collecting the required configurations (Fig. 1).

2.1.1. High-quality description at multiple applied bias values

A high-quality dataset and rigorous labelling are essential elements to provide an accurate description of electrified interfaces via ML-FFs. Any artefact in the solvent representation or ambiguities in the definition of the electrode potential can introduce inconsistencies in the labels, which become particularly detrimental when working with the small datasets employed in this work.

TRECI achieves a high-quality modelling of the interfaces via (a) a comprehensive description of the solvent dynamics, (b) an accurate treatment of the liquid water network, and (c) a robust electrification scheme (cf. left side of Fig. 1). The comprehensive description of the solvent is achieved by following a fully explicit solvent approach without any artificial vacuum or implicit solvent region. This allows for capturing in the dataset both the overall dielectric response of the solvent and the localised interactions between molecules and the surface.

A high-quality description of the liquid water network implies the choice of advanced DFT functionals due to the well-known shortcomings of the common PBE-GGA approximation in water-based systems [30,31]. We develop the datasets using the meta-GGA SCAN functional [32] as the DFT reference (cf. Section 4.1).

The inclusion of the effects of the applied bias in the framework of ML-FFs represents the most crucial element. This is realised within an electronically grand-canonical ensemble using the "double reference method" [28]. By employing this scheme, TRECI constructs multiple datasets corresponding to distinct potential energy surfaces (PESs), each associated with a well-defined and macroscopically accessible quantity, the bias value. The "double reference method" determines the electrode potential by using a series of potential references computed in auxiliary systems. This approach eliminates artefacts from solvent-vacuum interfaces and avoids inconsistencies between the energies computed in the systems with and without extra electrons (see Section 4.2). As a result, it provides a consistent relationship between surface charge and electrode potential, removes ambiguities in data labelling, and guarantees that each PES is uniquely identified. Such a rigorous electrification scheme represents the main step ahead with respect to other ML works.

To the best of our knowledge, a single framework integrating all these three modelling choices has not been realised in prior studies, largely due to the high computational cost associated with each component: SCAN is approximately 2.5 times more expensive than PBE, and a simulation with the "double reference method" is up to 10 times more costly than a calculation without external bias. When used together, these methods become prohibitively expensive, calling for highly data-efficient approaches.

2.1.2. ML force-fields via transfer learning

Given the high cost of generating such data across multiple potentials, training independent ML-FFs for each applied bias value would be highly inefficient. TRECI instead adopts a transfer learning paradigm [33] that reuses information from models trained in the absence of an electric field to efficiently learn the corresponding constant applied bias PESs.

Specifically, we draw on the recently introduced Franken framework [34]. In this scheme, energies and forces are predicted using a large-scale kernel regression model. Differently from the hand-crafted descriptors of traditional kernel-based ML-FFs, here they are built upon the node features of a pre-trained graph neural network (GNN), such as the general-purpose potentials from the MACE-MP family [35], inheriting the representations learnt by the deep neural network (cf. Section 4.3.1). Within TRECI, we extend this idea into a multi-head architecture, where a shared set of GNN-based descriptors is reused to simultaneously learn several PESs corresponding to different applied biases. Each head is trained to address a specific bias value (Fig. 2).

Because only a small regression block is bias-specific, the overall data requirement is drastically reduced. Furthermore, TRECI generalises the Franken concept beyond fine-tuning of universal ML-FFs: it also leverages domain-specific models trained for the same system at the PZC, in the absence of an applied field. In the following we refer to models constructed using general-purpose descriptors and domain-specific ones as Franken-MP0 and Franken-PZC, respectively. As detailed in Section 2.2.1, this combination of general-purpose and domain-specific models provides a route from stable exploratory simulations to accurate production runs.

2.1.3. *Data-efficient active learning*

The efficiency of TRECI ultimately hinges on identifying the few thermodynamically relevant configurations at each target bias. Simply re-labelling equilibrium structures obtained in the absence of an electric field is insufficient, as the applied potential can stabilise entirely different interfacial arrangements. To address this, TRECI adopts an iterative data acquisition strategy based on active learning. The dataset is progressively expanded by exploring the PESs with preliminary ML-FFs and selecting a limited number of new configurations for DFT labelling.

Specifically, for the selection, we exploit the capabilities of the "Data-Efficient Active Learning" (DEAL) protocol [36], which evaluates the similarity among the local environments of the sampled configurations to identify a reduced set of geometries to label. This filtering step is based on the Bayesian predictive variance of a Sparse Gaussian Process model (see Section 4.3.3). This approach yields a compact and non-redundant dataset while ensuring efficient coverage of the relevant configurational space.

2.1.4. *Complete strategy*

Here we summarise the steps of TRECI to construct ML-FFs at constant potential. The workflow begins with the construction of an ML-FF trained for the PZC case using the MACE architecture (see Section 4.3.2). The PZC model is developed to have domain-specific descriptors useful for the transfer learning. In addition, this gives a starting dataset for the

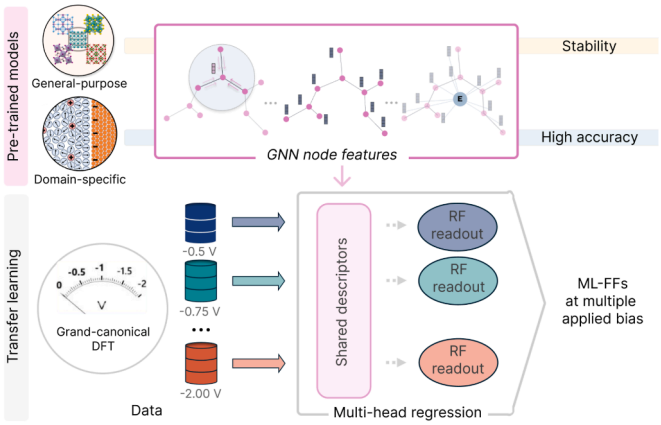

Fig. 2. TRECI strategy for the training of the constant potential ML-FFs via transfer learning. The node features of a pre-trained general-purpose or domain-specific GNNs are used as descriptors for a multi-head Franken model. Each readout block implements a large-scale kernel regression model targeting an applied bias value. Kernel functions are approximated by means of Random Fourier Features (RF).

ML-FFs with applied bias and serves as a baseline for the results in the absence of an electric field.

An iterative strategy composed of four steps is then applied (Fig. 1). The cycle starts with electronically grand-canonical DFT calculations. Using this data, ML-FFs are optimised for the corresponding constant bias values using transfer learning, employing both general-purpose (from the MACE-MP family) and domain-specific descriptors (from the model trained at the PZC). The models are then benchmarked for stability: if domain-specific descriptors fail to yield stable MD trajectories, the workflow proceeds using the general-purpose ones. Instead, once enough data are available and Franken-PZC models become robust, they are preferred due to their superior accuracy. Then, new atomic configurations are sampled using MD and selected through a data-efficient

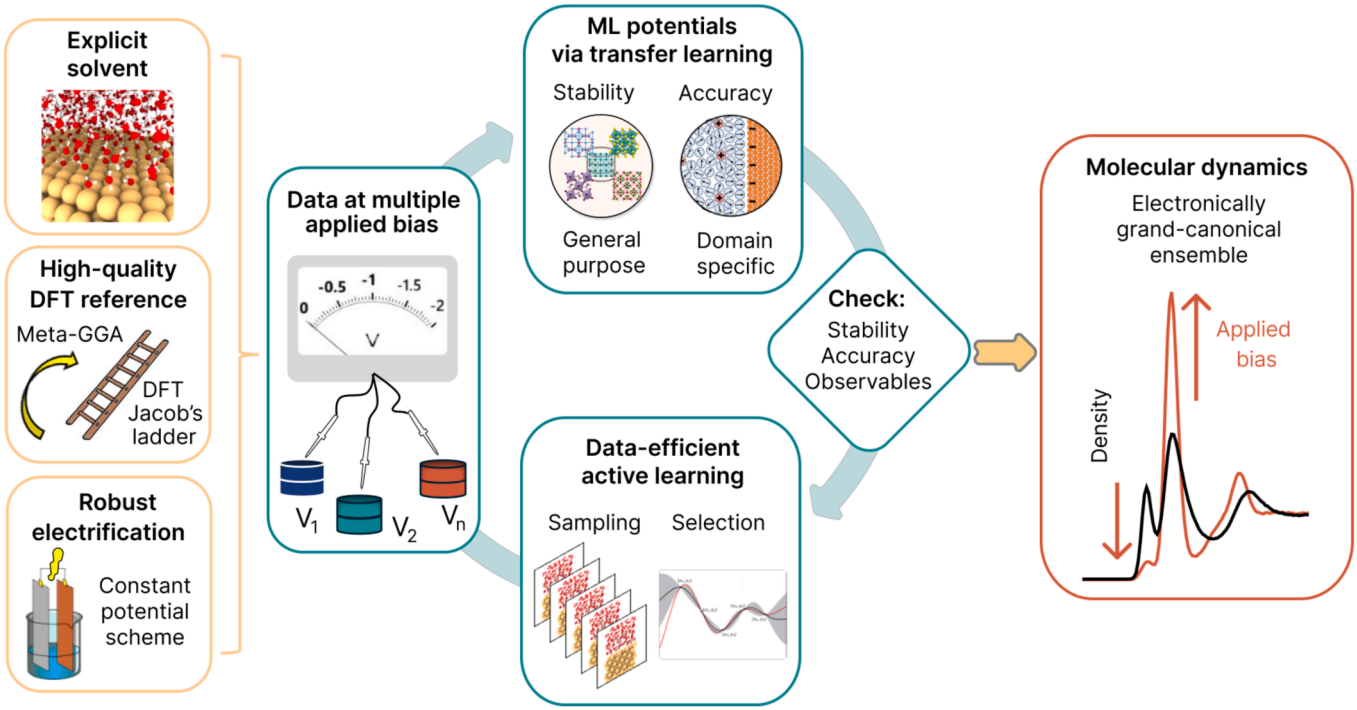

Fig. 1. General scheme of TRECI for electronically grand-canonical ML-FFs. The pillars of the methods are a high-quality dataset at multiple applied bias values, a data-efficient ML architecture via transfer learning and an effective active learning approach.

active learning protocol. The cycle continues until both the accuracy of energy and force predictions, as well as relevant physical observables, reach convergence. Only at this point, the ML-FFs are considered sufficiently reliable for use in production-level MD simulations.

2.2. Validation of the workflow

We validate the workflow of TRECI on a real case system. In particular, we study the copper–water interface at negative potentials, a representative cathodic system used in electrochemical applications such as the CO_2 electro-reduction reaction [7]. Fig. 3 shows the evolution of the learning process, from which it is clear that after a few active learning iterations and the collection of one thousand structures, TRECI achieves stable, accurate and converged results.

2.2.1. From stability to accuracy

We begin constructing the constant bias models using a subset of 250 configurations selected from the PZC dataset and labelled with grand-canonical DFT. This small dataset proves insufficient for training not only a standard MACE architecture but also domain-specific Franken-PZC ML-FFs: trajectories generated with these models lack stability. In contrast, the fine-tuned Franken-MP0 consistently yields stable results across all applied potentials and throughout all active learning cycles (Fig. 3a). For this reason, we start expanding the dataset in an active learning framework using configurations from MD trajectories generated with Franken-MP0 models. As the dataset is enlarged, Franken-PZC ML-FFs begin to produce stable results: at low applied bias ($V = -0.50$ V), stability is achieved after a single active learning cycle, while more negative potentials require additional iterations (Fig. 3a). In particular, we initially use equivariant domain-specific models, which are computationally more expensive but become robust in earlier iterations (Supplementary Section S1.2), and then we transition to invariant ML-FFs for production MD simulations.

Switching from general-purpose to domain-specific descriptors leads to a reduction in the prediction error (approximately 30% lower RMSE), ensuring a uniform accuracy across all the bias values (cf. Fig. 3b and Table 1). Notably, the accuracy achieved with PZC descriptors is comparable to the corresponding PZC MACE model one, but TRECI obtains these results using an amount of data equal to 25% of the PZC dataset.

Table 1
Root-Mean-Squared Error (RMSE) in meV/Å for the prediction of forces using TRECI models at different bias values, with descriptors from either MACE-MP0 (Franken-MP0) or our domain-specific PZC model (Franken-PZC).

Case [V vs. SHE]	−0.50	−0.75	−1.00	−1.25	−1.50	−1.75	−2.00
Franken-MP0	31.8	30.9	32.2	31.7	29.8	29.9	30.5
Franken-PZC	24.5	23.7	24.3	23.0	21.3	22.1	22.3

2.2.2. From accuracy to convergence

When using Franken-PZC models, the force accuracy rapidly converges to a plateau (Fig. 3b). Despite this convergence, we perform a more rigorous check by analysing the evolution of key physical observables. As a representative quantity, we monitor the peak heights in the oxygen density profile at the interface, which are highly sensitive to structural changes induced by the applied bias (see Section 2.3.1). The analysis of these quantities during the active learning provides insights into the learning process (Fig. 3c). As the models are progressively refined, the first peak height (chemisorbed molecules) decreases while the second one (physisorbed water) increases until both converge to stable values. This modulation of the peak density is a well-known effect of the applied potential, and these variations during the active learning cycles are an expected behaviour for a transfer learning process from an electric-field-agnostic model to ML-FFs in the presence of applied bias. By tracking these physical observables, we notice that even after the convergence of the force RMSE, approximately three additional active learning iterations are required for the peak densities to reach a plateau. At this point, we further validate the convergence of the solvent density

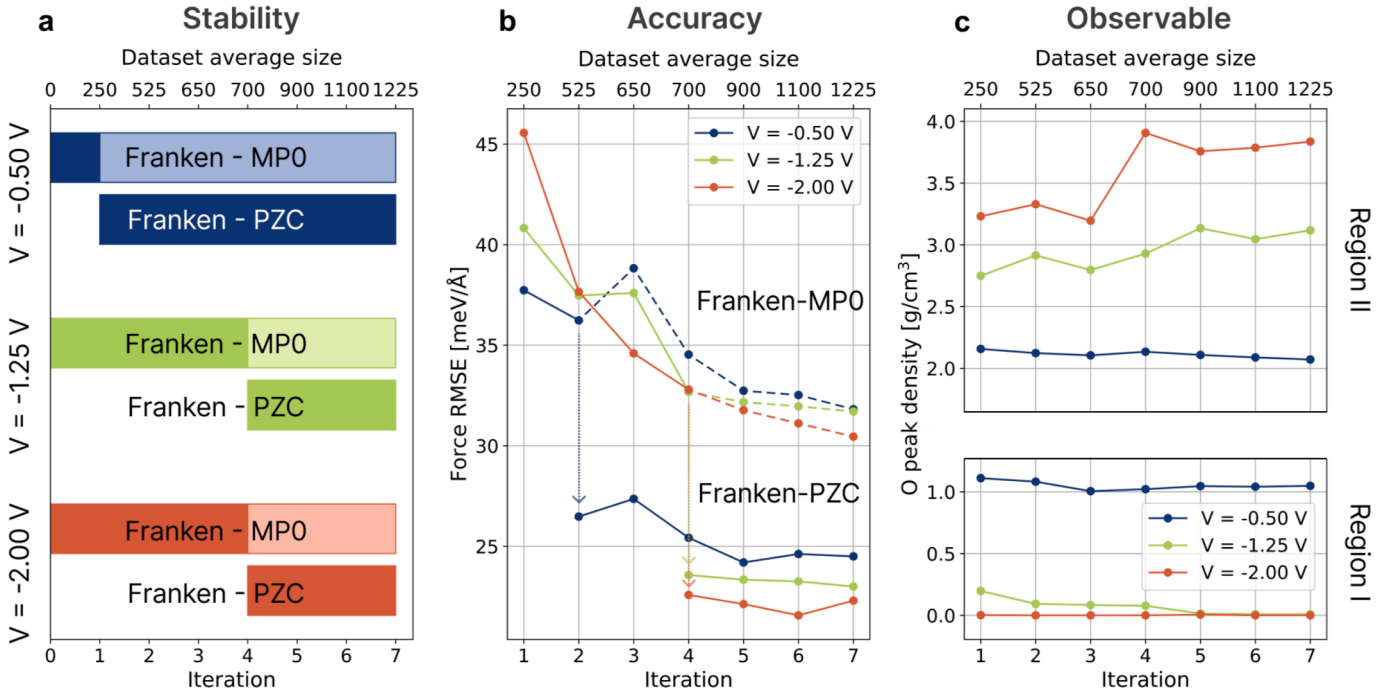

Fig. 3. Evolution of the learning process for TRECI models at representative target biases during the active learning. (a) Stability range of the different descriptors (Franken-MP0 vs. Franken-PZC): coloured bars show the iterations at which the ML-FFs are stable; dark colours identify the descriptors employed for the sampling at the different iterations. (b) Force RMSE for Franken-MP0 and Franken-PZC. Solid lines identify the descriptors employed for the sampling. The accuracy is not reported when the models are not stable. (c) Peak value of the oxygen density in the different interfacial regions.

profiles, comparing TRECI models built upon different descriptors, as detailed in Supplementary Section S1.3. Even if the starting models for the transfer learning are different, the ML-FFs predict reciprocally consistent observables, confirming the effectiveness of TRECI's strategy.

2.3. Application to Cu(111)–water interface

Having validated TRECI's workflow, the ML-FFs are employed in large-scale MD simulation to investigate the properties of interfacial water in contact with a Cu(111) surface, focusing on the effects of reducing bias. Based on an estimation of the PZC for this system (≈−0.1 to −0.2 V vs. SHE, see Section 4.2), we consider values smaller than −0.50 V up to −2.00 V vs. SHE as reducing potentials.

2.3.1. Structure of interfacial water

Firstly, we examine the density profile of interfacial water at different potentials (Fig. 4a). Looking at these profiles, it is possible to identify, at least, three visible regions in the solvent, close to the interface. These regions are sequentially labelled as I, II and III, moving from the surface towards the water bulk. There are still some minimal fluctuations in the density after 8 Å far from the surface, but they are completely smoothed out at 12 Å.

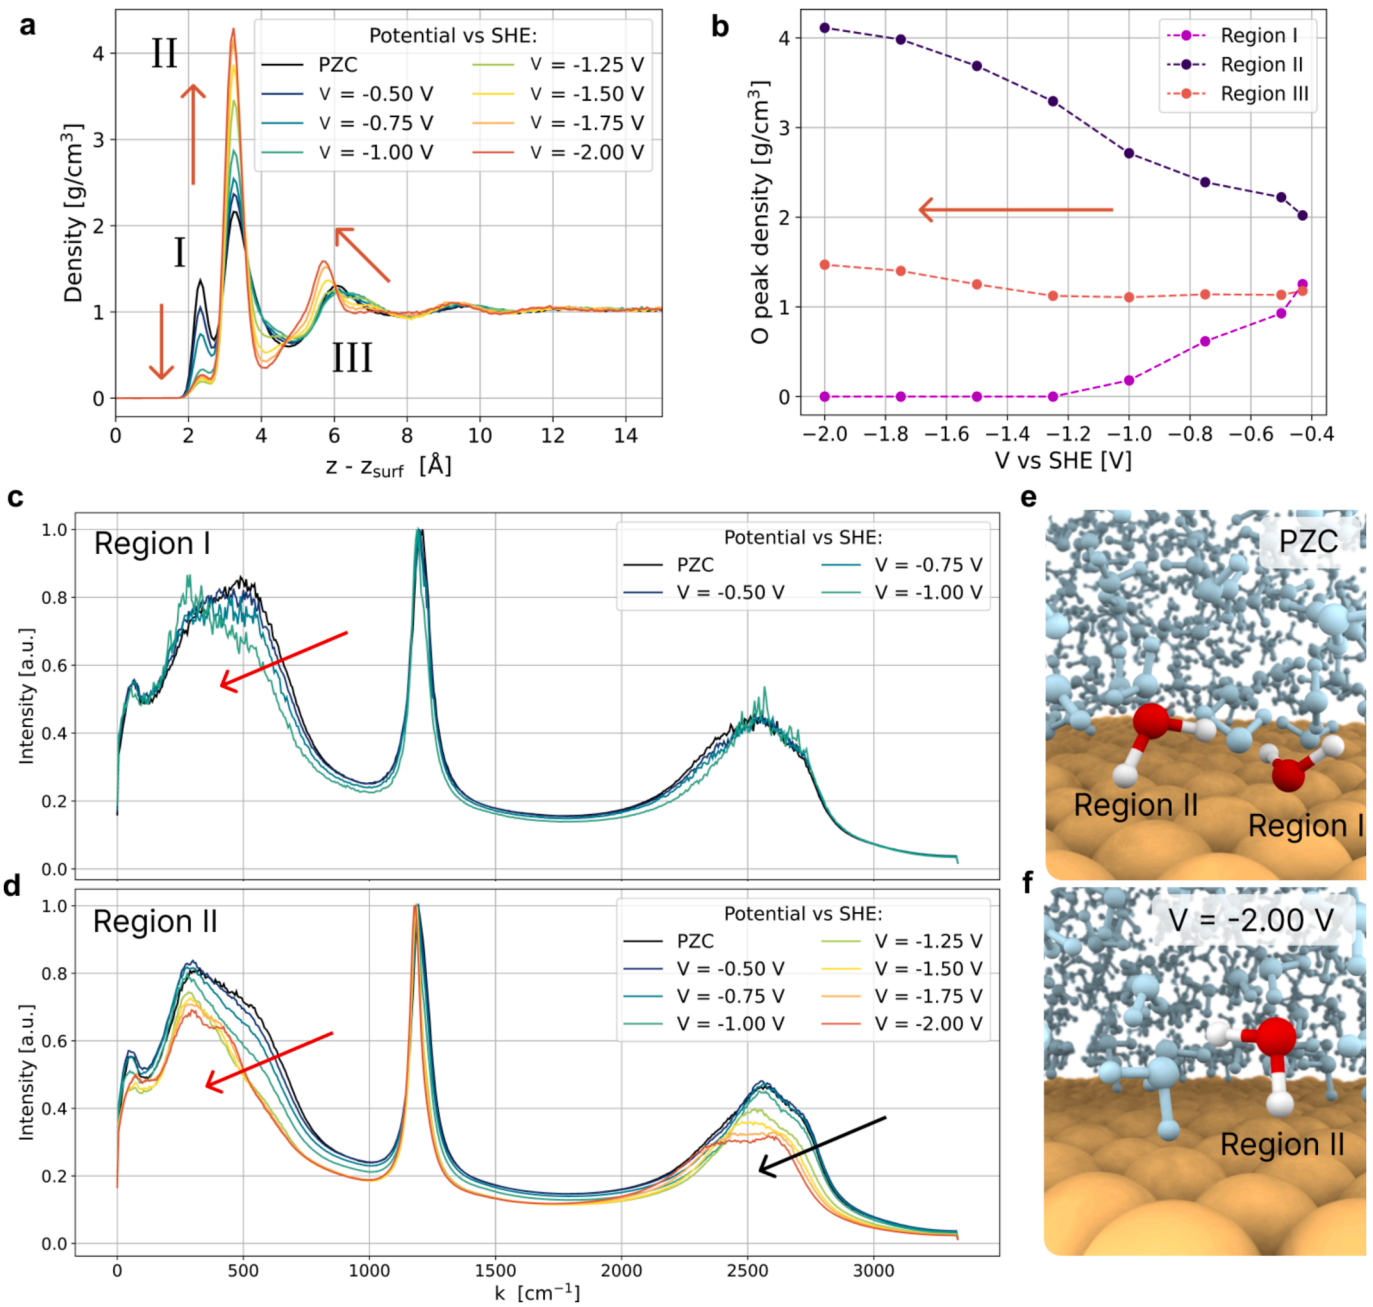

Fig. 4. Properties of interfacial water. (a) Density profile in the solvent region vs. the distance z with respect to the surface z_{surf}, at different values of external potential. Arrows emphasise trends moving towards more negative potentials. (b) Summary of the oxygen peak density height in the interfacial regions vs. the applied potential. VDOS spectra for hydrogen atoms in (c) region I and (d) region II. From low to high frequencies, it is possible to identify the peaks associated with the H-bond bending (≈40–50 cm^{-1}), the libration (150–700 cm^{-1}), the H–O–H bending (≈1200 cm^{-1}), the O–H symmetric and asymmetric stretching (2300–2800 cm^{-1}) [37,38]. Spectra are computed using the deuterium mass for the hydrogen species. Frames of the MD trajectory in which specific molecules are emphasised to show the typical water orientation in regions I and II (e) at PZC and (f) at −2.00 V vs. SHE. O and H atoms of emphasised molecules are depicted in red and white, respectively.

The most evident signature of the applied potential is the variation in the magnitude of the peak density (Fig. 4a). The density of peak I drastically decreases at increasing negative potentials, and it almost disappears for V ≤ -1.25 V. A more detailed analysis, focusing only on the oxygen role (Fig. 4b), reveals that the O contribution in region I is completely suppressed for V ≤ -1.25 V: there are no more H_2O molecules in region I, and the residual contribution to the total density in region I is associated with hydrogens of region II molecules.

The application of the bias also causes the density in peak II to increase to a value 4.5 times larger than the bulk water one (Fig. 4b). As discussed in the Supplementary Section S3, this peak modulation is associated with a variation in the preferential orientation of the water dipoles. Region I is composed of molecules with oxygen atoms exposed towards the surface and hydrogen atoms pointing away towards other water molecules (Fig. 4e). This orientation is compatible with a chemisorption process due to a dative bond between the oxygen lone pair and a copper atom. In contrast, molecules in region II are characterised by an orientation in which H atoms are mainly exposed towards the surface (Fig. 4e): these molecules are considered as physisorbed. This dipole flip mechanism with increasing negative bias has already been reported [39, 40] and is related to the variation of the strength of the Cu···O chemical interaction and the Cu···H electrostatic attraction at variable surface charge. At reducing potentials, the increase of the negative surface charge enhances the repulsion of negative oxygen atoms while attracting the positive hydrogens, flipping the water molecules (Fig. 4f). These effects are also illustrated in the Supplementary Video through MD snapshots, highlighting the Bader charge computed on the electrode at the DFT level and the behaviour of interfacial molecules.

The effectiveness of TRECI in studying multiple PESs is demonstrated by its ability to replicate these well-known effects and, more importantly, to produce monotonic and systematic trends among the various potential values. Even if AIMD can already give some information about these phenomena, our model provides a more detailed picture, affording an analysis at several applied potentials and without finite-size effects. In particular, two new elements are identified: an evident transition in region I and II between -1.00 and -1.25 V and a more complex behaviour of region III. As soon as peak I disappears, between -1.00 and -1.25 V, there is an abrupt increase in the O density in region II, sharpening the peak. At the same time, this transition also has consequences in region III. Here, the O density is almost unperturbed at low voltages (Fig. 4b), but peak III markedly increases at higher negative values. The bias influences not only the height of the peak but also its position: peak III tends to shift closer to the surface at more negative V ($\Delta z_{peak} \approx 0.3$ Å at $V = -2.00$ V vs. SHE). To the best of our knowledge, the effects of the applied potential in region III have never been discussed before. Indeed, this region and the complete transition to the water bulk are commonly not accessible via AIMD due to the computational cost of considering water layers thicker than 6–8 Å.

2.3.2. *Features of the interfacial H-bond network*

Another important feature of interfacial water is the strength of the H-bonds and how this is altered by the applied bias. This information can be inferred from the analysis of the vibrational density of states (VDOS) spectra for H atoms (Figs. 4c and d). For example, the relative strength of the H-bonds versus the O–H covalent bonds in water is correlated to O–H stretching vibration modes [37]. As emphasised with the black arrow in Fig. 4d, the O–H stretching band in region II visibly red-shifts. This shift is the typical signature of the weakening of the O–H covalent bond due to stronger H-bonds [37].

The stretching band reflects the strength of H-bonds through the relative motion of oxygen and hydrogen atoms along the covalent bond axis, while the libration modes capture the restricted rotational motion of water molecules caused by H-bonding interactions [41]. As emphasised by the red arrows in Figs. 4c and d, the libration band in regions I and II is altered by the applied bias. The shift at lower energies in region I is proof that molecules can more easily rotate in region I due to the weaker Cu···O interaction. Whereas, in region II, the high-energy shoulder of the libration band is more suppressed at increasing negative potentials. This red-shift can be interpreted as the signature of a less restricted rotational motion at increasing negative potentials due to a smaller number of H-bonds in region II (H atoms pointing towards the surface are not involved in H-bonds with other molecules).

It is important to note that computing VDOS spectra via AIMD is often limited by short simulation times, which hinder statistical convergence and hide spectral features. In contrast, the spectra presented here benefit from nanosecond-scale sampling, significantly improving the resolution of peak positions and revealing fine details in the spectra. This enhanced statistical quality allows for a more reliable interpretation of bias-induced changes in interfacial water dynamics.

3. Conclusions

We presented TRECI, a general and data-efficient workflow for developing ML-FFs to investigate metal–H_2O interfaces under external electrical bias. This enables a high-fidelity description of these systems by combining a fully explicit solvent model, a DFT-accurate representation of the liquid water network via the SCAN functional, and a constant potential framework that avoids explicit–implicit solvent compromises through the "double reference method". To address the computational cost of this rigorous setup, TRECI leverages a data-efficient transfer learning algorithm for the ML-FFs training and an effective active learning protocol. To ensure model stability in the early stages of the ML-FF development, we perform transfer learning from general-purpose potentials. In subsequent iterations, we refine the ML-FFs by transferring information from domain-specific models to achieve a superior accuracy. This dual strategy enables the construction of robust and precise ML-FFs across a wide range of potential values. Our workflow achieves stable MD trajectories with only a few hundred configurations, reaches the same accuracy as models trained from scratch, and converges physical observables with a final dataset of just one thousand configurations. This data efficiency is critical for enabling DFT-level labelling in constant potential simulations, especially with our computationally demanding settings. Importantly, the applicability of the TRECI workflow is not constrained by our specific *ab initio* setup used for data labelling. Its modular design allows for easy customisation (in the choice of DFT functional, electrification scheme, or pre-trained model for transfer learning), making it adaptable to the requirements of diverse electrochemical interfaces, not limiting to metal–H_2O systems.

To demonstrate its capabilities, we applied the TRECI workflow to Cu–water interfaces as a representative system. Our ML models successfully revealed detailed structural and dynamical properties of electrified metal–water interfaces, achieving a level of resolution rarely accessible with conventional approaches. This picture can be further enriched by incorporating surface defects, ions, adsorbates, and electrocatalytic reactions to simulate specific operando conditions. These more complex scenarios, and in particular the explicit inclusion of the role of ions, will introduce new challenges. Charged species necessitate a more rigorous and systematic treatment of electrostatic interactions via long-range schemes in the ML models [42]. At the same time a faithful representation of realistic ionic concentrations requires a careful choice of large simulation cells and, possibly, specific algorithms that enable the realisation of a grand canonical ensemble with respect to ions [43]. These are enhancements of the TRECI capability that can be considered depending on the specific systems under study and the scientific questions one wants to address. TRECI is concerned as an open and modular framework that can accommodate such methodological extensions as system complexity increases, while preserving the fundamental advantages of its data-efficient and robust workflow for the development of more advanced ML-FFs. For these reasons, we believe that the TRECI strategy paves the way for high-quality, atomistic simulations of

electrified interfaces, driving progress in both the development of modelling techniques and the advancement of fundamental and applied research in electrochemistry.

4. Methods

4.1. Details of DFT calculations

The *ab initio* calculations required to generate data for the ML dataset are based on non-polarised DFT simulations, as implemented in the *Vienna Ab initio Simulation Package* (VASP), version 6.4.3 [44–46]. The projector-augmented-wave (PAW) method is employed to describe the electron-ion interaction [47]. Electronic wave functions are expanded in plane waves with a cut-off energy of 800 eV in conjunction with a dense FFT grid to properly guarantee numerical convergence of energies as well as of forces. The Brillouin zone is sampled by employing a Gamma-centred (11 11 1) Monkhorst-Pack mesh in the case of the unit cell of the Cu slab and consistently reduced for the supercells. A Gaussian smearing with a spreading of 0.1 eV is employed. Gamma-point-only calculations are performed for systems including only H_2O molecules.

The meta-generalised-gradient approximation (meta-GGA) "Strongly Constrained and Appropriately Normed" (SCAN) functional is employed [30,32]. Indeed, it is well known that common GGA functionals, such as PBE, tend to markedly over-structure the liquid water network [30,48]. Whereas, SCAN succeeds in describing the H_2O behaviour due to its capability to reproduce the correct relative magnitude of the covalent O–H bonds, the hydrogen bonds and the dispersion forces that determine the H-bond network [30]. In view of the numerical instability and convergence issues associated with the original formulation of the SCAN functional, the revisited r^2SCAN version is adopted in this work [49].

4.2. Interface electrification scheme

The electrification of the interface at the DFT level is performed following the "double reference method" [28]. By allowing the surface charge to adjust to reach a target potential and compensating with a homogeneous background charge for the added or removed electrons, this approach is compatible with an electronically grand-canonical ensemble [50]. The robust definition of the electrode potential is established via a chain of potential references involving two auxiliary systems such that this definition is not influenced by the counter-charge and any fictitious solvent-vacuum interface. Specifically, the first auxiliary system has no extra charge (i.e., no applied potential) and a vacuum region to define the vacuum level; the second is with no extra charge and no vacuum and represents a reference for the same system but with extra charge (i.e., with applied potential).

Here, we employ the definition of the electrode potential within the "double reference method" to accomplish two tasks. First, by applying this scheme to the configurations in the PZC dataset, we estimate the potential of zero charge to be approximately −0.1 to −0.2 V vs. SHE. This value serves as an upper bound for the range of reducing potentials. Second, by using the same method on geometries with added or removed charge, we establish a direct link between the net charge and the corresponding applied bias. This enables us to define a criterion for tuning the system charge to reach a target bias value. In this approach, the same geometry is labelled at multiple bias values. These calculations share the chain of references. As a result, the computational cost of the auxiliary systems is distributed across datasets corresponding to different applied biases.

Practically, the grand-canonical DFT is performed by modifying the "Fully Constant Potential" (FCP) ASE calculator proposed in Ref. [51], with the definition of the electrode potential in the "double reference method". This calculator is then integrated into a Python routine that implements an automatic workflow to evaluate the auxiliary systems of the "double reference method" and perform constant potential calculations at fixed geometry for more target potential values (see Code Availability).

4.3. Machine learning force-fields

4.3.1. Constant potential models via Franken transfer learning

The training of the constant potential force-fields is performed within the Franken transfer learning approach [34]. This scheme is based on the integration of the accurate descriptors of a pre-trained deep neural network with an efficient large-scale kernel regression approach. The kernel function is approximated by means of Random Fourier Features (RF) maps to overcome the scalability issues of kernel methods in the case of large datasets [52]. This structure allows for fast training and inference typical of kernel methods, but irrespective of the dataset size.

As shown in Fig. 2, starting from common descriptors (the backbone) extracted as the node features of a graph neural network (GNN), different readout blocks are used to target specific bias values: we trained seven RF-based readouts for potentials ranging from −0.50 V to −2.00 V vs. SHE (with a spacing of 0.25 V). This approach is inherently scalable, enabling the simultaneous generation of a greater number of regression blocks that encompass an even broader spectrum of applied biases.

The peculiarity of the transfer learning in TRECI is the use of multiple backbones in the different cycles of active learning: our strategy relies on transfer learning from general-purpose models during the initial stages of active learning, followed by the use of domain-specific descriptors at later stages and in MD production runs. The MACE-MP0 family of universal models is selected as the general-purpose backbone because it is trained on a dataset comprising 89 chemical species and has been explicitly validated for its ability to qualitatively reproduce the properties of a broad range of solid-water interfaces by extrapolating reasonably well beyond its training domain [35]. This makes it a general choice not only for Cu-water systems but also for other interfacial systems. At the same time, a growing number of MACE universal models are now available, increasing the flexibility in the choice of the general-purpose descriptors without compromising the generality of TRECI's methodology: general-purpose models serve primarily as backbones to ensure stability rather than to maximise accuracy. The latter is achieved by using domain-specific backbones. These are extracted from a MACE model trained on a set of data of the same interface but without an applied electric field (cf. Section 4.3.2). While this approach requires training an additional model, the computational cost of labelling configurations at PZC is significantly lower (approximately 10–20 times less than performing full constant potential calculations across seven target bias values). Moreover, the computational cost associated with developing an ML model from scratch is substantially reduced by applying the DEAL protocol also in this context. The superior expressiveness and data efficiency of the MACE architecture relative to other commonly employed ML schemes further reduce the requirements on the size of the PZC dataset [53]. The result is a data-efficient strategy also in the generation of a PZC model from scratch (cf. Section 4.3.2). In view of these considerations, the increased accuracy with a reduced number of expensive grand-canonical DFT evaluations justifies the additional, but moderate, effort required to develop the PZC model.

After testing different descriptors and numbers of RFs (Supplementary Section S1.1), final results are achieved with the descriptors from our PZC MACE model and 8192 RFs.

4.3.2. MACE model for the PZC

The ML force-fields (ML-FFs) at the PZC are built using the MACE software, version 0.3 [54]. The peculiarity of MACE is the construction of a high-order message-passing scheme based on a many-body atomic cluster expansion (ACE) representation and the adoption of equivariant internal features for the message [55]. It results in a data-efficient structure with appreciable in-domain and out-domain predictions. Specifically, we employ a 4-body term expansion and two interaction layers, each with a cut-off radius of 6 Å. Considering the possible presence of non-local

contributions outside this receptive field, we test models with cut-off radii up to 9 Å, as well as a model with a cut-off of 5 Å but with three interaction layers (Supplementary Section S2). In all cases, we observe that the RMSE of the force predictions decreases only marginally as the receptive field increases. At the same time, the solvent density profile remains unchanged with the expansion of the receptive field, indicating that non-local contributions do not significantly affect the solvent structure. In addition, Supplementary Section S2 provides a more stringent analysis of the length scale over which interactions remain relevant in the systems considered here. Our results show that, in our electrified metal-water interfaces, these interactions are effectively screened by the surrounding water beyond approximately 8–10 Å. This screening substantially suppresses the influence of contributions of interactions outside the receptive field of our ML-FFs. These results demonstrate that the selected cut-off radius and number of interaction layers are sufficient to capture all relevant interactions within the system.

During the construction of the ML-FF, an architecture with equivariant messages ($L = 1$) and 128 channels is employed to better optimise the stability of the model in case of a limited-size database. Once the dataset is complete, the final ML-FF is trained with invariant messages ($L = 0$) and 256 channels, resulting in a reasonable balance between accuracy and computational efficiency. The dataset is split into training/validation/test subsets with a ratio of 85:10:5. The model is optimised with the AMSGrad algorithm, using a learning rate of 0.01, a batch size of 4 and a maximum number of epochs equal to 1000. Performance is evaluated on energy and forces with a weighted Root-Mean-Squared-Error (RMSE) loss function. The weights for energy and forces in the loss function are initially set to 1 and 100, respectively. In the last 20% of the training, the weight for energy is increased to 1000.

4.3.3. Dataset generation

Different strategies are adopted to generate a preliminary dataset for the PZC. A pre-existing dataset [56] is used to describe the pure water bulk. Configurations for bare Cu surfaces are included after sampling AIMD trajectories at a lower accuracy level (i.e., PBE). A ML universal model [35] is also used to collect data for Cu–H_2O systems. For all these cases, the configurations are labelled consistently with our DFT setup.

Starting from this initial pool of configurations, the training dataset is iteratively refined through the "Data-Efficient Active Learning" (DEAL) scheme, recently introduced in Ref. [36]. This method allows us to progressively enlarge the dataset and refine the ML-FF in a data-efficient manner, providing new configurations selected from MD trajectories obtained with preliminary ML-FFs. Firstly, the MD configurations are screened, relying on the uncertainty prediction obtained by a *query-by-committee* approach. To avoid redundant labelling, this set of selected configurations is then filtered by DEAL using a sparse Gaussian process model, trained on-the-fly on the pre-selected data. More details on this approach are reported in the original article [36]. This strategy ensures broad and uniform coverage of the relevant configuration space while dramatically reducing the number of costly *ab initio* calculations.

At the end of the active learning, the complete dataset for the Cu–H_2O interface at the PZC is composed of about 5320 configurations, including systems of bulk H_2O, bare Cu surfaces and Cu–H_2O interfaces.

A similar iterative strategy is used to construct datasets at multiple applied potentials. Starting from a subset of 250 configurations sampled from the PZC dataset, models are progressively refined using the DEAL protocol. In this case, a two-step DEAL selection is performed. First, configurations are filtered independently for each applied bias; then, the selected configurations are merged across all bias values and screened again alongside the existing dataset. This second DEAL selection reduces redundancy among configurations sampled at the different biases and with respect to the existing dataset. This iterative refinement continues until the TRECI models reach an accuracy comparable to the reference PZC model and demonstrate stable convergence of key physical observables.

4.4. Molecular dynamics simulations

Classical molecular dynamics (MD) simulations are performed with the Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS) software [57], supplemented by MACE v0.3 [54]. NVT simulations are performed with an integration step of 0.5 fs, and the deuterium mass is used for the hydrogen atoms to properly describe the dynamics of light atoms. The temperature is controlled using a Nosé–Hoover chain thermostat with a damping parameter of 100 times the integration step. The thermostat temperature is properly set to ensure that the outermost metal layers of the slab and water molecules are in thermodynamic equilibrium at 330 K. This slight increase above room temperature is required to overcome the residual over-structuring of water predicted by SCAN [30]. The dynamical properties are computed in the microcanonical ensemble to remove any spurious contribution from the thermostat. For the VDOS spectra, the velocity autocorrelation function is evaluated by recording the velocities every 5 fs for several time windows of 6 ps. MD runs are performed in a slab system composed of six layers of metal atoms, with the two innermost ones fixed. Periodic boundary conditions are imposed in all directions, and no regions of vacuum are present. MD runs are performed in a slab system with a lateral size of 10–15 Å (4 × 4 up to 6 × 6 supercells) during the active learning phase and 20 Å (8 × 8 supercells) in the final production run. In the latter case, a region of 30 Å (i.e., 15 Å per side) is filled with water molecules. Previous studies have demonstrated that the system dimensions employed in the final production runs are sufficiently large for both structural and dynamical properties to be independent of the system size [19]. The complete recovery of a constant water density at distances of 11–12 Å far from the surface, without residual oscillations in the density profile for each value of applied potential, provides additional evidence that the newly observed features in the properties of the aqueous solvent reported here do not arise from finite-size artefacts.

CRediT authorship contribution statement

Michele Giovanni Bianchi: Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Investigation, Formal analysis. **Michele Re Fiorentin:** Writing – review & editing, Formal analysis. **Francesca Risplendi:** Writing – review & editing, Formal analysis. **Candido Fabrizio Pirri:** Writing – review & editing, Formal analysis. **Michele Parrinello:** Writing – review & editing, Formal analysis. **Luigi Bonati:** Writing – review & editing, Writing – original draft, Visualization, Supervision, Methodology, Formal analysis, Conceptualization. **Giancarlo Cicero:** Writing – review & editing, Supervision, Formal analysis, Conceptualization.

Data availability

A minimal dataset to evaluate the TRECI workflow is provided alongside the code at https://github.com/michelegiovannibianchi/TRECI. The molecular dynamics trajectories underlying the analysis of the interface are available on Zenodo at [58].

Code availability

The TRECI code underlying this work is freely available at https://github.com/michelegiovannibianchi/TRECI. The FCP calculator employed in TRECI is available at https://github.com/michelegiovannibianchi/DoubleReferenceMethod-FCP-calculator. TRECI workflow requires the MACE code available at https://github.com/ACEsuit/mace, the Franken code available at https://github.com/CSML-IIT-UCL/franken and the DEAL code available at https://github.com/luigibonati/DEAL.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors acknowledge the support of the Data Science and Computation Facility at the Fondazione Istituto Italiano di Tecnologia and the CINECA award under the ISCRA initiative. M.G.B. is grateful to F. Raffone, P. J. Buigues, and S. Perego for valuable discussions, and to P. Novelli and G. Meanti for providing the Franken code. L.B. and M.P. acknowledge funding from the European Union - NextGenerationEU initiative and the Italian National Recovery and Resilience Plan (PNRR) from the Ministry of University and Research (MUR), under Project PE0000013 CUP J53C22003010006 “Future Artificial Intelligence Research (FAIR)”. The authors are also grateful for the financial support from the National Recovery and Resilience Plan (NRRP), Mission 4 “Education and Research” – Component 2 “From research to business” - Investment 3.1 “Fund for the realization of an integrated system of research and innovation infrastructures” – Call n. 3264 of 28/12/2021 of Italian Ministry of University and Research | Award Decree n. 128 (21/06/2022)- Project code: IR0000027, Concession Decree No. 128 of 21/06/2022 adopted by the Italian Ministry of Research, CUP: B33C22000710006, Project title: iENTRANCE. This work also received support by the European Innovation Council (EIC) under the European Union’s Horizon Europe research and innovation programme through the project PREDICT (Grant Agreement No. 101219185). The views and opinions expressed are those of the authors and do not necessarily reflect those of the European Union or the European Innovation Council. Neither the European Union nor the granting authority can be held responsible for them. Finally, this publication is supported by the COST Action CA24103: EUropean COllaborative Network on electroCatalysis for Efficient Renewable Technologies (EU-CONCERT) and the flagship program "Teaching Science to Computers" of the Italian Institute of Technology.



Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.esci.2026.100642.

References

[1] J. Carrasco, A. Hodgson, A. Michaelides, A molecular perspective of water at metal interfaces, Nat. Mater. 11 (2012) 667–674.
[2] G. Gonella, E.H.G. Backus, Y. Nagata, D.J. Bonthuis, P. Loche, A. Schlaich, R. R. Netz, A. Kühnle, I.T. McCrum, M.T.M. Koper, M. Wolf, B. Winter, G. Meijer, R. K. Campen, M. Bonn, Water at charged interfaces, Nat. Rev. Chem 5 (2021) 466–485.
[3] D. Kolb, Reconstruction phenomena at metal-electrolyte interfaces, Prog. Surf. Sci. 51 (1996) 109–173.
[4] X. Xia, M.L. Berkowitz, Electric-field induced restructuring of water at a platinum-water interface: a molecular dynamics computer simulation, Phys. Rev. Lett. 74 (1995) 3193–3196.
[5] T. Iwasita, F. Nart, *In situ* infrared spectroscopy at electrochemical interfaces, Prog. Surf. Sci. 55 (1997) 271–340.
[6] G. Cicero, A. Calzolari, S. Corni, A. Catellani, Anomalous wetting layer at the Au (111) surface, J. Phys. Chem. Lett. 2 (2011) 2582–2586.
[7] S. Nitopi, E. Bertheussen, S.B. Scott, X. Liu, A.K. Engstfeld, S. Horch, B. Seger, I.E. L. Stephens, K. Chan, C. Hahn, J.K. Nørskov, T.F. Jaramillo, I. Chorkendorff, Progress and perspectives of electrochemical CO_2 reduction on copper in aqueous electrolyte, Chem. Rev. 119 (2019) 7610–7672.
[8] A. Auer, X. Ding, A.S. Bandarenka, J. Kunze-Liebhäuser, The potential of zero charge and the electrochemical interface structure of Cu(111) in alkaline solutions, J. Phys. Chem. C 125 (2021) 5020–5028.
[9] J.-X. Zhu, J. Cheng, Machine learning potential for electrochemical interfaces with hybrid representation of dielectric response, Phys. Rev. Lett. 135 (2025) 018003.
[10] A. Groß, S. Sakong, Ab initio simulations of water/metal interfaces, Chem. Rev. 122 (2022) 10746–10776.
[11] O.T. Unke, S. Chmiela, H.E. Sauceda, M. Gastegger, I. Poltavsky, K.T. Schütt, A. Tkatchenko, K.-R. Müller, Machine learning force fields, Chem. Rev. 121 (2021) 10142–10186.
[12] J.-B. Le, X.-H. Yang, Y.-B. Zhuang, M. Jia, J. Cheng, Recent progress toward ab initio modeling of electrocatalysis, J. Phys. Chem. Lett. 12 (2021) 8924–8931.
[13] N.G. Hörmann, S.D. Beinlich, K. Reuter, Converging divergent paths: constant charge vs constant potential energetics in computational electrochemistry, J. Phys. Chem. C 128 (2024) 5524–5531.
[14] G. Kastlunger, P. Lindgren, A.A. Peterson, Controlled-potential simulation of elementary electrochemical reactions: proton discharge on metal surfaces, J. Phys. Chem. C 122 (2018) 12771–12781.
[15] S.M.R. Islam, F. Khezeli, S. Ringe, C. Plaisance, An implicit electrolyte model for plane wave density functional theory exhibiting nonlinear response and a nonlocal cavity definition, J. Chem. Phys. 159 (2023) 234117.
[16] D. Le, An explicit-implicit Hybrid Solvent Model for Grand Canonical Simulations of the Electrochemical Environment, 2023, https://doi.org/10.26434/chemrxiv-2023-z2n4n.
[17] T.W. Ko, J.A. Finkler, S. Goedecker, J. Behler, A fourth-generation high-dimensional neural network potential with accurate electrostatics including non-local charge transfer, Nat. Commun. 12 (2021) 398.
[18] C. Zhang, M.F. Calegari Andrade, Z.K. Goldsmith, A.S. Raman, Y. Li, P.M. Piaggi, X. Wu, R. Car, A. Selloni, Molecular-scale insights into the electrical double layer at oxide-electrolyte interfaces, Nat. Commun. 15 (2024) 10270.
[19] S.K. Natarajan, J. Behler, Neural network molecular dynamics simulations of solid–liquid interfaces: water at low-index copper surfaces, Phys. Chem. Chem. Phys. 18 (2016) 28704–28725.
[20] A.E.G. Mikkelsen, J. Schiøtz, T. Vegge, K.W. Jacobsen, Is the water/Pt(111) interface ordered at room temperature? J. Chem. Phys. 155 (2021) 224701.
[21] P.S. Rice, Z.-P. Liu, P. Hu, Hydrogen coupling on platinum using artificial neural network potentials and DFT, J. Phys. Chem. Lett. 12 (2021) 10637–10645.
[22] X. Chen, M. El Khatib, P. Lindgren, A. Willard, A.J. Medford, A.A. Peterson, Atomistic learning in the electronically grand-canonical ensemble, npj Comput. Mater. 9 (2023) 73.
[23] N. Bergmann, N. Bonnet, N. Marzari, K. Reuter, N.G. Hörmann, Machine learning the energetics of electrified solid-liquid interfaces, Phys. Rev. Lett. 135 (2025) 146201.
[24] X. Tian, A. Tosello Gardini, U. Raucci, H. Xiao, Y. Zhuo, M. Parrinello, Electrochemical potential-driven water dynamics control CO2 electroreduction at the Ag/H2O interface, Nat. Commun. 16 (2025) 10636.
[25] L. Chen, Y. Tian, X. Hu, S. Chen, H. Wang, X. Zhang, Z. Zhou, Constant-potential MD with neural network potentials reveals cation effects on CO2 reduction at Au-water interfaces, JACS Au 6 (2025) 868–880.
[26] M.S. Chen, J. Lee, H.-Z. Ye, T.C. Berkelbach, D.R. Reichman, T.E. Markland, Data-efficient machine learning potentials from transfer learning of periodic correlated electronic structure methods: liquid water at AFQMC, CCSD, and CCSD(T) accuracy, J. Chem. Theor. Comput. 19 (2023) 4510–4519.
[27] V. Zaverkin, D. Holzmüller, L. Bonfirraro, J. Kästner, Transfer learning for chemically accurate interatomic neural network potentials, Phys. Chem. Chem. Phys. 25 (2023) 5383–5396.
[28] C.D. Taylor, S.A. Wasileski, J.-S. Filhol, M. Neurock, First principles reaction modeling of the electrochemical interface: consideration and calculation of a tunable surface potential from atomic and electronic structure, Phys. Rev. B 73 (2006) 165402.
[29] M. Naher, M.A. Gonzálvez, C.M. Williams, P.V. Bernhardt, Emerging applications and mechanistic insights of copper mediated electrocatalysts in organic transformations, Chem. Soc. Rev. 54 (2025) 7304–7338.
[30] M. Chen, H.-Y. Ko, R.C. Remsing, M.F. Calegari Andrade, B. Santra, Z. Sun, A. Selloni, R. Car, M.L. Klein, J.P. Perdew, X. Wu, Ab initio theory and modeling of water, Proc. Natl. Acad. Sci. U.S.A. 114 (2017) 10846–10851.
[31] T.E. Gartner, L. Zhang, P.M. Piaggi, R. Car, A.Z. Panagiotopoulos, P.G. Debenedetti, Signatures of a liquid–liquid transition in an ab initio deep neural network model for water, Proc. Natl. Acad. Sci. U.S.A. 117 (2020) 26040–26046.
[32] J. Sun, A. Ruzsinszky, J. Perdew, Strongly constrained and appropriately normed semilocal density functional, Phys. Rev. Lett. 115 (2015) 036402.
[33] J. Falk, L. Bonati, P. Novelli, M. Parrinello, M. Pontil, Transfer learning for atomistic simulations using GNNs and kernel mean embeddings, arXiv:2306.01589 (2024).
[34] P. Novelli, G. Meanti, P.J. Buigues, L. Rosasco, M. Parrinello, M. Pontil, L. Bonati, Fast and Fourier features for transfer learning of interatomic potentials, npj Comput. Mater. 11 (2025) 293.
[35] I. Batatia, P. Benner, Y. Chiang, A.M. Elena, D.P. Kovács, J. Riebesell, X. R. Advincula, M. Asta, M. Avaylon, W.J. Baldwin, F. Berger, N. Bernstein, A. Bhowmik, F. Bigi, S.M. Blau, V. Cărare, M. Ceriotti, S. Chong, J.P. Darby, S. De,

F. Della Pia, V.L. Deringer, R. Elijošius, Z. El-Machachi, E. Fako, F. Falcioni, A. C. Ferrari, J.L.A. Gardner, M.J. Gawkowski, A. Genreith-Schriever, J. George, R.E. A. Goodall, J. Grandel, C.P. Grey, P. Grigorev, S. Han, W. Handley, H.H. Heenen, K. Hermansson, C.H. Ho, S. Hofmann, C. Holm, J. Jaafar, K.S. Jakob, H. Jung, V. Kapil, A.D. Kaplan, N. Karimitari, J.R. Kermode, P. Kourtis, N. Kroupa, J. Kullgren, M.C. Kuner, D. Kuryla, G. Liepuoniute, C. Lin, J.T. Margraf, I.-B. Magdău, A. Michaelides, J.H. Moore, A.A. Naik, S.P. Niblett, S.W. Norwood, N. O'Neill, C. Ortner, K.A. Persson, K. Reuter, A.S. Rosen, L.A.M. Rosset, L. L. Schaaf, C. Schran, B.X. Shi, E. Sivonxay, T.K. Stenczel, C. Sutton, V. Svahn, T. D. Swinburne, J. Tilly, C. van der Oord, S. Vargas, E. Varga-Umbrich, T. Vegge, M. Vondrák, Y. Wang, W.C. Witt, T. Wolf, F. Zills, G. Csányi, A foundation model for atomistic materials chemistry, J. Chem. Phys. 163 (2025) 184110.

[36] S. Perego, L. Bonati, Data efficient machine learning potentials for modeling catalytic reactivity via active learning and enhanced sampling, npj Comput. Mater. 10 (2024) 291.

[37] J. Le, Q. Fan, L. Perez-Martinez, A. Cuesta, J. Cheng, Theoretical insight into the vibrational spectra of metal–water interfaces from density functional theory based molecular dynamics, Phys. Chem. Chem. Phys. 20 (2018) 11554–11558.

[38] S. Jin, X. Fan, C. Stamper, R.A. Mole, Y. Yu, L. Hong, D. Yu, M. Baggioli, On the temperature dependence of the density of states of liquids at low energies, Sci. Rep. 14 (2024) 18805.

[39] M.F. Toney, J.N. Howard, J. Richer, G.L. Borges, J.G. Gordon, O.R. Melroy, D. G. Wiesler, D. Yee, L.B. Sorensen, Distribution of water molecules at Ag(111)/electrolyte interface as studied with surface X-ray scattering, Surf. Sci. 335 (1995) 326–332.

[40] F. Raffone, R. Khatib, M. Sulpizi, C. Cucinotta, Revealing the molecular interplay of coverage, wettability, and capacitive response at the Pt(111)-water solution interface under bias, Commun. Chem. 8 (2025) 58.

[41] Y. Tong, T. Kampfrath, R.K. Campen, Experimentally probing the libration of interfacial water: the rotational potential of water is stiffer at the air/water interface than in bulk liquid, Phys. Chem. Chem. Phys. 18 (2016) 18424–18430.

[42] F. Grasselli, K. Rossi, S. de Gironcoli, A. Grisafi, Long-range electrostatics in atomistic machine learning: a physical perspective, arXiv:2602.11071 (2026).

[43] A. Soyemi, K. Baral, T. Szilvási, Modeling equilibrium solid–liquid interfaces under effective constant chemical potential using machine learning interatomic potentials, J. Phys. Chem. A 129 (2025) 11245–11255.

[44] G. Kresse, J. Hafner, Ab initio molecular-dynamics simulation of the liquid-metal–amorphous-semiconductor transition in germanium, Phys. Rev. B 49 (1994) 14251–14269.

[45] G. Kresse, J. Furthmüller, Efficiency of ab-initio total energy calculations for metals and semiconductors using a plane-wave basis set, Comput. Mater. Sci. 6 (1996) 15–50.

[46] G. Kresse, J. Furthmüller, Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set, Phys. Rev. B 54 (1996) 11169–11186.

[47] G. Kresse, D. Joubert, From ultrasoft pseudopotentials to the projector augmented-wave method, Phys. Rev. B 59 (1999) 1758–1775.

[48] M.D. LaCount, F. Gygi, Ensemble first-principles molecular dynamics simulations of water using the SCAN meta-GGA density functional, J. Chem. Phys. 151 (2019) 164101.

[49] J.W. Furness, A.D. Kaplan, J. Ning, J.P. Perdew, J. Sun, Accurate and numerically efficient r^2SCAN meta-generalized gradient approximation, J. Phys. Chem. Lett. 11 (2020) 8208–8215.

[50] A. Hagopian, M.-L. Doublet, J.-S. Filhol, T. Binninger, Advancement of the homogeneous background method for the computational simulation of electrochemical interfaces, J. Chem. Theor. Comput. 18 (2022) 1883–1893.

[51] Z. Xia, H. Xiao, Grand canonical ensemble modeling of electrochemical interfaces made simple, J. Chem. Theor. Comput. 19 (2023) 5168–5175.

[52] G. Dhaliwal, P.B. Nair, C.V. Singh, Machine learned interatomic potentials using random features, npj Comput. Mater. 8 (2022) 7.

[53] D.P. Kovács, I. Batatia, E.S. Arany, G. Csányi, Evaluation of the MACE force field architecture: from medicinal chemistry to materials science, J. Chem. Phys. 159 (2023) 044118.

[54] I. Batatia, D.P. Kovács, G.N.C. Simm, C. Ortner, G. Csányi, MACE: Higher order equivariant message passing neural networks for fast and accurate force fields, arXiv:2206.07697 (2023).

[55] I. Batatia, S. Batzner, D.P. Kovács, A. Musaelian, G.N.C. Simm, R. Drautz, C. Ortner, B. Kozinsky, G. Csányi, The design space of E(3)-equivariant atom-centred interatomic potentials, Nat. Mach. Intell. 7 (2025) 56–67.

[56] T.I. Gartner, L. Zhang, P. Piaggi, R. Car, A. Panagiotopoulos, P. Debenedetti, Data from "Signatures of a liquid-liquid Transition in an Ab Initio Deep Neural Network Model for Water", 2020, https://doi.org/10.34770/45m3-am91.

[57] A.P. Thompson, H.M. Aktulga, R. Berger, D.S. Bolintineanu, W.M. Brown, P. S. Crozier, P.J. in't Veld, A. Kohlmeyer, S.G. Moore, T.D. Nguyen, R. Shan, M. J. Stevens, J. Tranchida, C. Trott, S.J. Plimpton, LAMMPS - a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales, Comput. Phys. Commun. 271 (2022) 108171.

[58] M.G. Bianchi, M. Re Fiorentin, F. Risplendi, C.F. Pirri, M. Parrinello, L. Bonati, G. Cicero, Electrochemical Interfaces at Constant Potential: Data-Efficient Transfer Learning for machine-learning-based Molecular Dynamics, 2025, https://doi.org/10.5281/zenodo.17664081.