Generalised deep-learning workflow for the prediction of hydration layers over surfaces

Yashasvi S. Ranawat [a], Ygor M. Jaques [a], Adam S. Foster [a,b]

[a] *Department of Applied Physics, Aalto University, P.O. Box 11100, 00076 Aalto, Espoo, Finland*
[b] *WPI Nano Life Science Institute (WPI-NanoLSI), Kanazawa University, Kakuma-machi, Kanazawa 920-1192, Japan*

ARTICLE INFO

ABSTRACT

Atomic force microscopy (AFM) is paving the way for understanding the solid–liquid interfaces at the nanoscale. These AFM studies are complemented with molecular dynamics (MD) simulations of hydration layers over candidate surfaces for a comprehensive characterisation. We earlier proposed, in Ranawat *et.al.* (2021), a deep-learning (DL) network to predict hydration layers over the candidate surfaces much more rapidly than computationally-intensive MD. However, the proposed elements-as-channels based network is bound to the elements present in the training surfaces. Here, we develop a generalised descriptor of the surface to train element-agnostic networks. We demonstrate the descriptor's efficacy by predicting the hydration layers over a dolomite surface using a network trained on the calcite and magnesite surfaces. We also demonstrate the transfer-learning capability of such a descriptor by incorporating mica into the training surfaces, and predict the pyrophyllite and boehmite surfaces. Further, we propose an energy-based DL framework to gauge the possible prediction accuracy of a network on surfaces hitherto unseen. We combine these advance techniques into a generalised work-flow to complement AFM studies.

1. Introduction

The nature of the complex hydration structures formed at solid–liquid interfaces plays a key role in many macroscopic surface phenomena [1,2], that drive various natural and technological processes [3]. Amongst many complementary techniques to characterise solid–liquid interfaces at the nanoscale, Atomic force microscopy (AFM) [4] has emerged as a leading method to image systems in solution at high speed with molecular resolution [5–14]. However, the convolution of tip-surface-solvent interactions and entropic effects in the measured force at high resolution means that it remains a challenge to provide comprehensive understanding and interpretation of the results - requiring extensive free energy simulations [15–18].

This challenge was bypassed, to some degree, by the introduction of methods that linked experimental AFM force data directly to water densities [19,20]. These approaches have been recently expanded to include the influence of the hydration features around the tip [21] and radius of the tip [17]. Ultimately, even these approaches necessitate systematic molecular dynamics (MD) simulations of water on the surface, which are accessible only to

experts and only if good classical force field parameterisations are available. These challenges render the high resolution AFM imaging in liquids remains limited in application, especially considering the rapid progress seen in molecular characterisation offered by functional-tip AFM at low-temperature and ultra-high vacuum [22,23,]. This is particularly troubling when we consider the breakthroughs seen in the application of AFM to biomolecular systems [24].

In our previous work we proposed a deep-learning (DL) U-Net network with elements-as-channels input [25]. This network demonstrated a rapid, robust and reliable prediction of hydration layers over surfaces, and defects, of polymorphs of calcite, directly from their atomic structures. However, binding the network inputs to certain elements (calcium, carbon, and oxygen in that case) posed a big disadvantage when surfaces with other elements were studied. In the present work, we introduce a novel four-channel descriptor to depict the surface structure, based on the MD force-field terms, namely, their van der Waals and Coulombic interactions (see Fig. 1). This element-agnostic design is shown to be advantageous in predicting various surfaces with a wide range of elemental composition. Further, this design is shown to improve the performance of the network on newer surfaces, through transfer learning [26], which pass on the knowledge trained from related surfaces by progressively augmenting the training set.
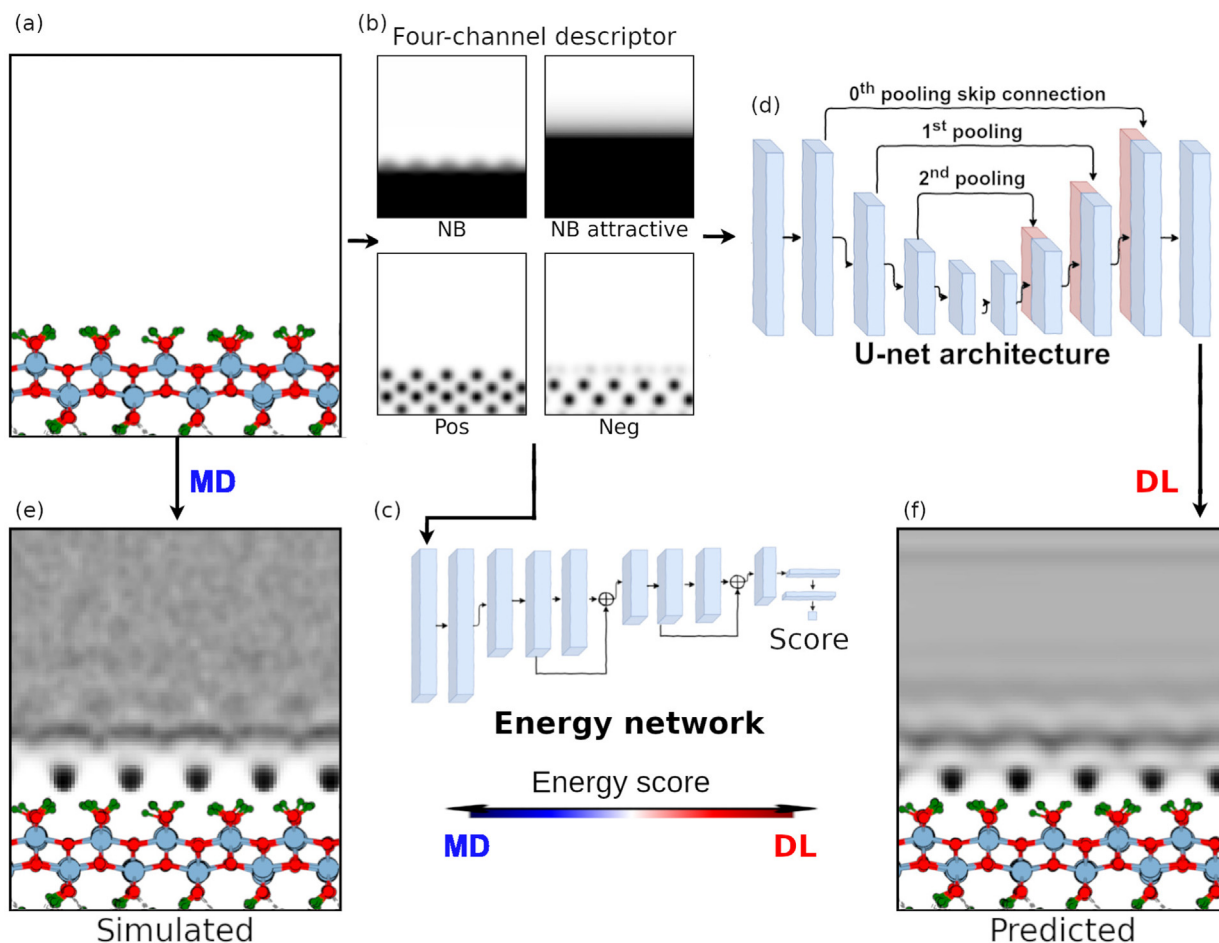
Fig. 1. Schematic of the workflow for hydration layer prediction over a surface (here, (010) boehmite with aluminium, hydrogen and oxygen atoms represented by grey, green, and red circles, respectively): (a) The atomic structure of the surface. (b) The four-channel density descriptor of the surface, comprising non-bonded interactions (NB), zoomed-in attractive region of the non-bonded interactions (NB attractive), the positive and the negative parts of the electrostatic interactions (Pos and Neg, respectively). (c) The energy network. (d) The deep U-net neural network, adapted from [25] (CC-BY). (e) Simulated hydration layers over the surface. (f) Predicted hydration layers over the surface using the DL network.

Notwithstanding the advantages, the proposed DL networks are *overconfident* in their predictions due to their deterministic nature [28]. Given a surface input, seen/unseen by the network during training, the hydration density are predicted with 100% certainty, even when the prediction is erroneous. To solve this issue, a secondary model, called an energy-based model, can provide out-of-distribution detection [29]. We adapted this energy-based DL framework to predict an "energy" or a score of the four-channel descriptor of the surface. This score is a robust figure-of-merit to gauge the representation (log-likelihood) of a new surface in the given training set. We thereby demonstrate a framework which estimates the possible accuracy with which network can predict the hydration layers over the given surface.

Our approach links strongly to the developing DL methods applied to data challenges in Scanning Probe Microscopy (SPM) [30–33]. In particular, the success of deep learning Convolutional Neural Networks (CNN) [34] in image recognition tasks has led to their application to the analysis of SPM images [35], especially in the context of molecular/surface feature/defect characterisation [36–41], scanning-probe characterisation [42,43], and techniques for autonomously-driven SPM [44–46].

However, no earlier studies, beyond ours, have applied DL to SPM at solid–liquid interfaces [47]. Furthermore, this work is an entirely different workflow approach from the past studies of DL applied to SPM, due to the system-agnostic descriptor, out-of-distribution detection technique, training data and DL model.

In the following sections, we present the complete workflow for the generalised tools to predict hydration layers over surfaces. This workflow comprises three parts. Firstly, a DL network that can rapidly predict hydration layers and screen candidate surfaces to characterise surface structures with hydration layers imaged from AFM, cif. Fig. 1. Secondly, a transfer-learning based training surface augmentation, afforded by element-agnostic surface descriptor design, to reliably include new surfaces and enhance the network capability. Lastly, an energy-based secondary DL model that benchmarks the possible accuracy of a network, given a set of training surfaces. Cumulatively, these methods aim to provide rapid, robust, and reliable AFM image analysis of hydration layers over varied surfaces.

2. Methods

2.1. MD and training data

We simulated $(10\bar{1}4)$ surfaces of calcite, magnesite, and dolomite using the inter-atomic potentials derived by Raiteri et al. [48], an extensively used forcefield for AFM imaging of carbonate minerals in liquids [16,14,49]. We chose the SPC/Fw flexible model for water [50]. In the case for (001) mica-water interactions, we used the CLAYFF forcefield [51] with TIP3P [52] water model, as used in mica-solvation studies in the context of AFM imaging

[53]. We performed the MD simulations using the LAMMPS MD code [54].

These forcefields are functionally made of bonded (bonds, angles, and dihedrals) and non-bonded interactions. In the used forcefields, the non-bonded interactions, primarily responsible for surface-water interactions, comprised pairwise and electrostatic interactions, given respectively as:

$$
\begin{aligned}
E_{pairwiseLennard-Jones} &= 4\epsilon\left(\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6}\right) \\
E_{pairwiseBuckingham} &= Ae^{-r/\rho} - \frac{C}{r^6}
\end{aligned}
\tag{1}
$$

$$
E_{electrostatic} = \frac{q_1 q_2}{4\pi\epsilon_o r^2}
\tag{2}
$$

Hence, we designed a four-channel surface descriptor that encapsulates these non-bonding interactions. The pairwise interactions were taken from the pairwise interaction potentials between the surface atoms and water molecule, which were either Lennard-Jones or Buckingham type potential, cif. Eq. 1. To better capture the repulsive and the attractive parts of the non-bonded forcefield in a normalised input — required for effective network training — these fields were split into separate channels. Secondly, the electrostatic interactions were represented by the product of surface atom's charge and the charge on oxygen of the water, cif. Eq. 2. This accommodated the different water models used across training surfaces. The product-charges were smeared with a gaussian atomic density to get an equivalent electrostatic forcefield. Similar to the non-bonded forcefield, the electrostatic forcefield was also split into the positive and the negative electrostatic forcefield.

As in our previous approach [25], we created a large dataset of surface defects by removing permutations of the top surface atoms in the calcite, the magnesite, and the mica surface (see Supplementary Information for more details). We generated 39,708 cases, 27,504 cases, and 49,872 cases of the calcite, the magnesite, and the mica surfaces, respectively. These cases included a, $10 \times 10 \times 20$ Å^3 volume, four-channel descriptor of the surface for input and a target of hydration layer density over the surface, cif. Fig. 1 (b, e). Finally, these cases were distributed into training, validation, and test sets. The training set was used to optimise the network parameters. The validation set was used to determine the progress of training the network and avoid over-fitting. Lastly, the test set was used to gauge the unbiased performance of the trained network. In this work, an approximate split of 70%, 20%, and 10% of the dataset was used for training, validation, and testing, respectively. Given the large dataset in this work, a smaller test set was considered to be sufficient [55]. Moreover, the work also includes tests on newer surfaces with reasonable results, which show consistent trends in test errors making us confident with the split used.

2.2. Machine learning

For training, we used the DL U-Net architecture, comprising three pooling scales and skip connections [56], as described in our previous work [25], cif. Fig. 1 (d). However, the model was re-purposed to take-in inputs of 4 channels. Although the network had to be trained from scratch, the element-agnostic approach of the four-channel descriptor lead to simpler integration of surfaces with varied surface elements. Newer surfaces were accommodated without the need to re-train a new network with random initial weights. We first trained our network on the (10$\bar{1}$4) surfaces of calcite. Then we progressively augmented the training data set, one at a time, with the (10$\bar{1}$4) magnesite and (001) mica surfaces, to demonstrate the transfer-learning aspect of the network. Further, we predicted increasingly challenging surfaces, that were not represented in the training sets of the networks, to gauge the generalisability of the network with progressive data augmentation.

For out-of-distribution detection, an energy-based DL framework [57,29] was designed that follows the encoder part of the U-Net architecture, cif. Fig. 1 (c). The network takes in the four-channel descriptor as input, and predicts a scalar in the form of a score associated with the log-likelihood of the input with respect to the training data. A contrastive divergence scheme [58] was used during training of the network, which maximised log-likelihood (or the energy score) for in-distribution surfaces while minimising that for out-of-distribution surfaces (see Supplementary Information for more details). In other words, a higher energy score from the network is better as it indicates that a surface is well represented in the training set — this is a convention for energy networks in machine learning and opposite to the intuitive understanding of energy from physics.

3. Results

We used the hydration layer densities over (10$\bar{1}$4) dolomite, (010) boehmite, (010) pyrophyllite surfaces in addition to the test cases of (10$\bar{1}$4) calcite, magnesite, and (001) mica to benchmark the four-channel descriptor, see Fig. 2. We trained the network on an extensive set of the calcite surfaces and progressively add magnesite and mica surfaces to monitor the training. We compare the simulated and predicted densities by plotting the 1D mean density along the *z* direction for the dolomite surface. We also plot the 2D mean density, and *xy* slices of the hydration layers along the *z* direction. Further, we plot the prediction errors to gauge the transfer learning ability of the element-agnostic approach, along with the score of the energy-based DL model, with the progressive inclusion of several surfaces in the training set. We use matplotlib [59] to plot the graphs.

In the first stage of training, the DL U-Net model was trained on the calcite and magnesite surfaces. The training data didn't contain any surface with dolomite (which has both calcium and magnesium atoms, Fig. 2). Hence, the network wasn't trained to understand the interplay of the hydration layers over the magnesium and the calcium atoms. Regardless, similar to the previous network[25], the network predicts the location of the hydration layers well, albeit with slightly offset magnitudes, cif. Fig. 3 (c, d). On comparing the 2D slices in Fig. 3 (a), the hydration density peaks clearly affect the neighbouring peaks. Moreover, the predicted hydration-density-peak features over the calcium and the magnesite layers are similar to that in the simulated case. The network is distinctly able to exhibit the difference in the hydration layers over the calcium and the magnesium atoms[60], and deduces the interplay between the corresponding hydration peaks reasonably. This justifies the choice of the descriptor, especially given that the magnesium and calcium elements differ very slightly in their pair-coefficient forcefield terms with oxygen in the water model, and have the same charge in the forcefield parameterised by Raiteri et al. [48].

In the second stage of training, three versions of the model were trained with cascading training sets. The first network was trained on the calcite surfaces, while the second network continued with the parameters of the first network, albeit trained on calcite and magnesite surfaces. Similarly, the third network used the parameters of the second network, and was trained on calcite, magnesite, and mica surfaces. The mean absolute error (MAE) — L1 error — of the three networks in predicting the hydration layers of the dolomite, pyrophyllite, and boehmite surfaces, and the test sets of the calcite, magnesite and mica surfaces, is plotted in Fig. 5 (a) of the three networks. It is clear that the progressive inclusion of varied surfaces in the training sets results in a decrease in the prediction error.
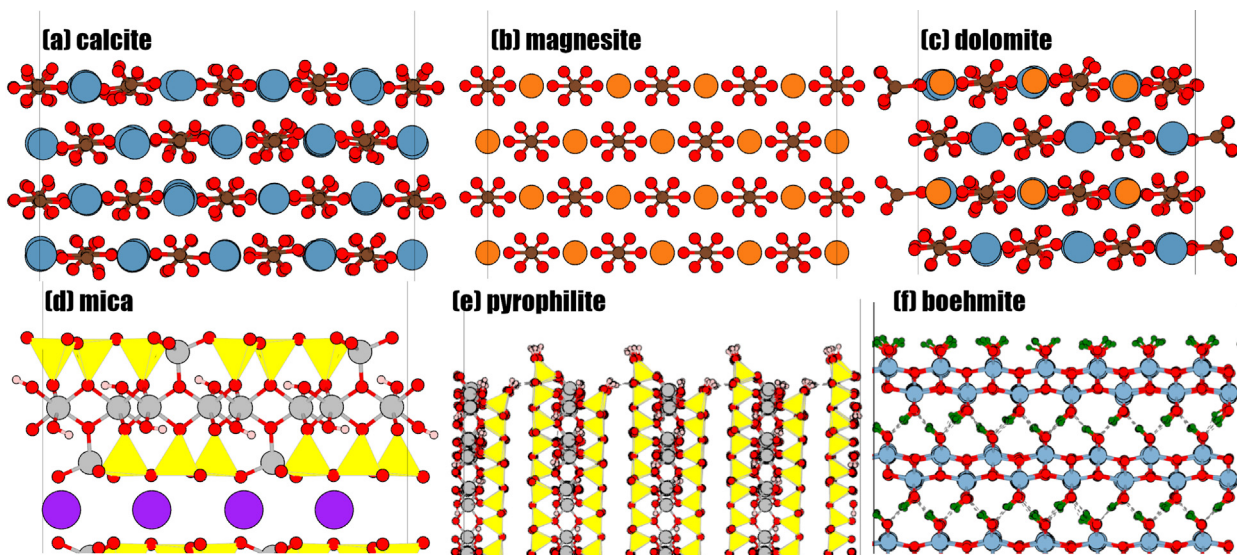
Fig. 2. Surface slab of (a) $(10\bar{1}4)$ calcite, (b) $(10\bar{1}4)$ magnesite, (c) $(10\bar{1}4)$ dolomite, (d) (001) mica, (e) (010) pyrophyllite, (f) (010) boehmite. The vertical lines represent the unit cell. The tetrahedral silicon-oxide structure is represented by a yellow tetrahedron. The atoms of carbon, oxygen, calcium, magnesium, aluminium, hydrogen, and potassium are represented with balls of brown, red, blue, orange, grey, green, and purple colour. The visualisations are made in VESTA [27].
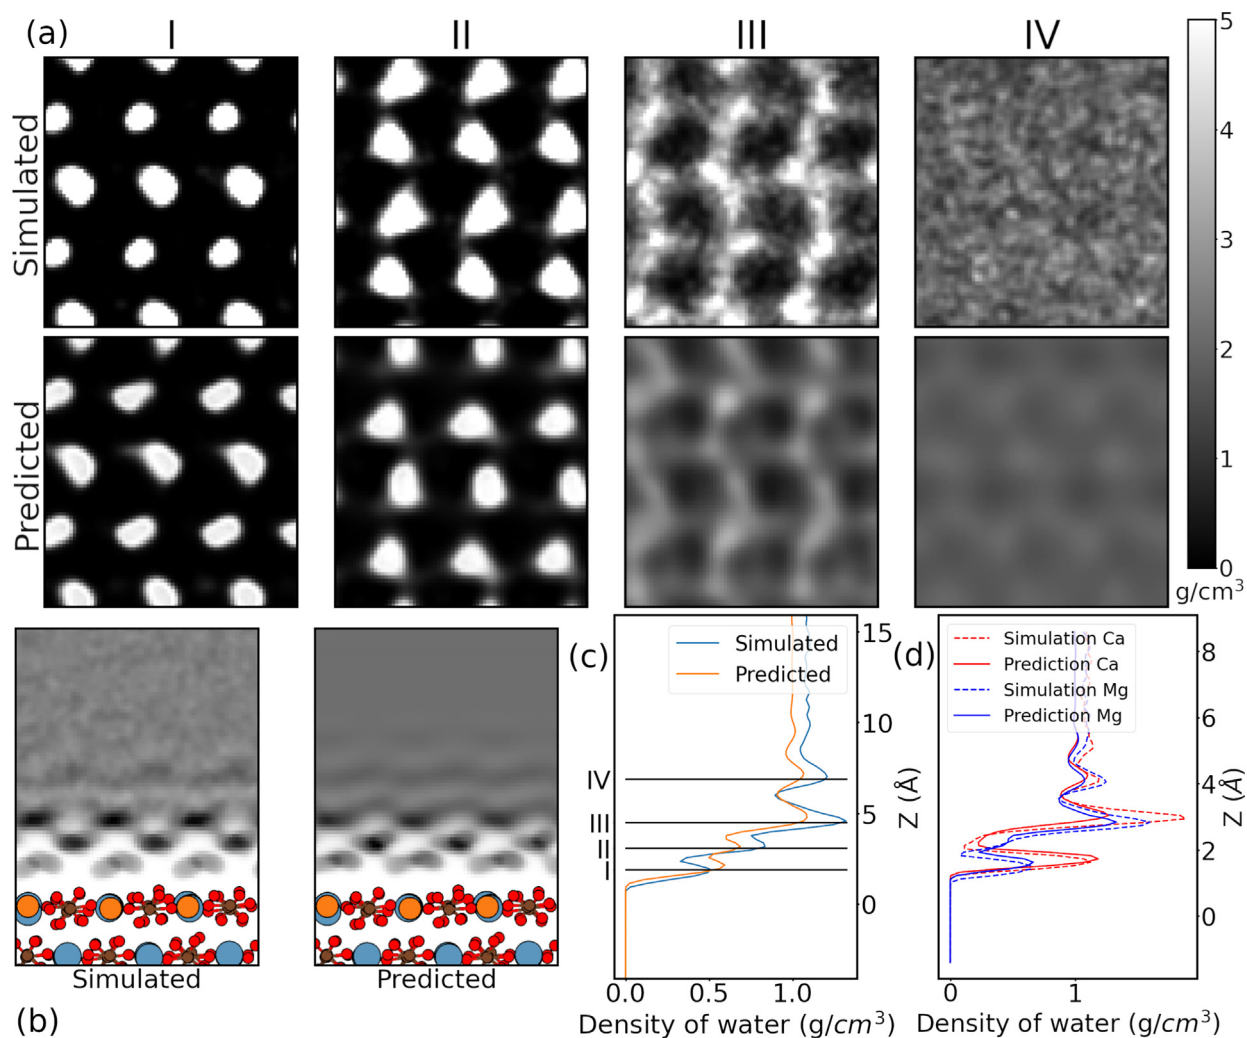

Fig. 3. Hydration layers prediction over $(10\bar{1}4)$ dolomite surface using the U-Net trained on $(10\bar{1}4)$ magnesite and calcite surfaces. (a) 2D slice comparison of the hydration layers in simulated and predicted hydration layer density. (b) 2D xz plane mean hydration layer density, (c) 1D z direction hydration layer density. (d) 1D hydration layers, along z axis, averaged over the calcium and the magnesium surface atoms. The blue, orange, red, and brown balls represent calcium, magnesium, oxygen, and carbon atoms, respectively.

To test the generality of the trained networks, hydration layers were predicted on the surfaces of (010) pyrophyllite and (010) boehmite. The MAE of the predictions are higher, cif Fig. 5 (a). We attribute this to the fact that the surface structure of Al-O-H, present in pyrophyllite and boehmite, is not well represented in the mica-surface training set. Furthermore, the hydration-layer peaks predicted have a slight offset, cif Fig. 4 (d, h). For the hydration layers over boehmite surface the features in the xy slice of the first peak are similar, as seen in Fig. 4 (e), while the DL network predicts faint patterned features in the second and third peak which are non-existent in the simulated images. In the case of pyrophyllite, Fig. 4 (a), the network highlights some features in the xy slices more than others, when compared to the simulated images. Moreover, in the slice corresponding to the second peak, the network fills more water in the void seen in the simulated

images, while capturing the light feature in the middle of the void. However, we note that our z-slice comparison is quite thorough and such differences will not be noticed in AFM images where the amplitude of the tip-oscillation convolves several z-slices [16].

In the final stage of the training, similar to the last step, three versions of the energy-based DL models were trained with cascading training sets. The energy values, or *score*, are plotted for the dolomite, pyrophyllite, and boehmite surfaces, along with the test sets of the calcite, magnesite and mica surfaces, using the three networks, as seen in Fig. 5 (b). A similar trend is observed where with progressive inclusion of varied surfaces in the training sets the predicted energy or the score of newer surfaces increases. However, due to the contrastive divergence method, the magnitude of the energy values are not normalised. Therefore, an external observable is needed to understand the energy values. Using
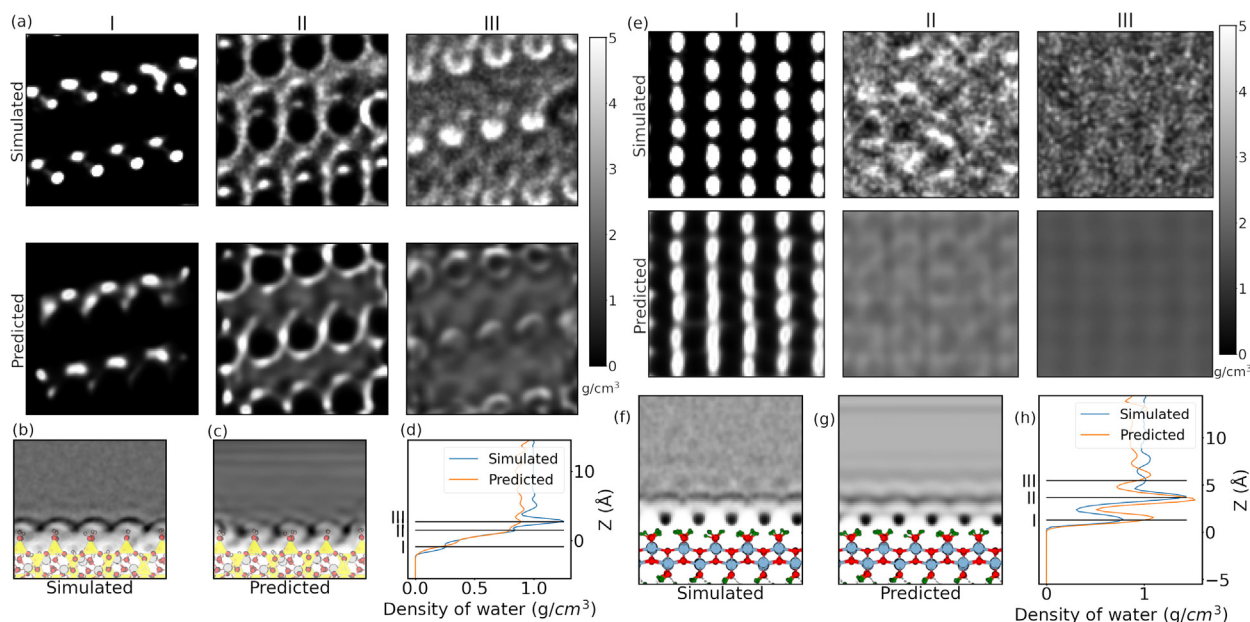

Fig. 4. Prediction of hydration layers over (010) pyrophyllite and (010) boehmite surface using the U-Net DL network trained on (001) mica, (10$\bar{1}$4) magnesite and calcite surfaces. (a)[(e)] 2D slice comparison of the hydration layers in simulated and predicted hydration layer density over pyrophyllite [boehmite]. (b, c)[(f, g)] simulated and predicted 2D xz plane mean hydration layer density over pyrophyllite [boehmite], (d)[(h)] 1D z direction hydration layer density over pyrophyllite [boehmite].
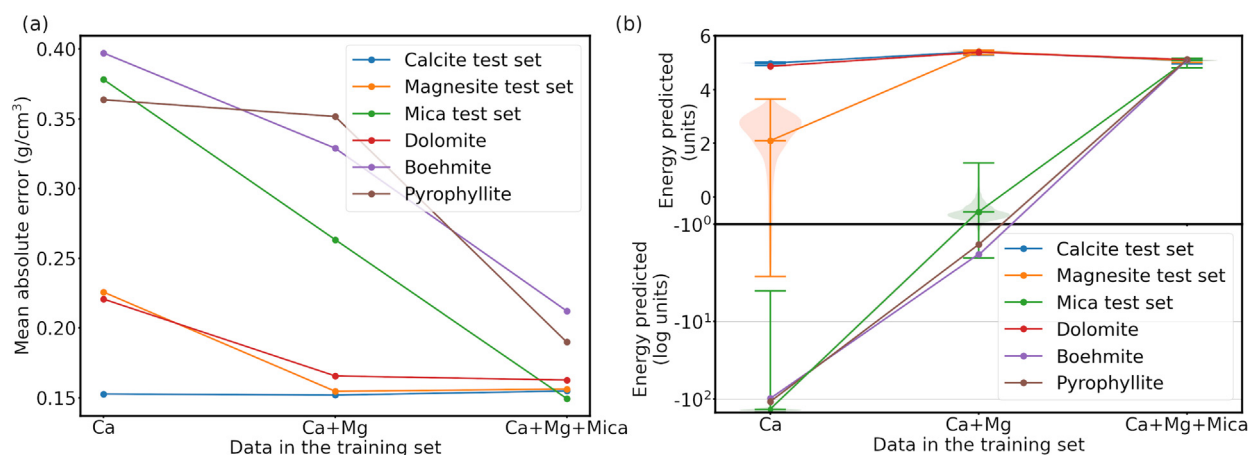

Fig. 5. (a) Mean absolute error in the prediction of hydration layers with networks trained on different training surfaces. (b) Energy values of surfaces with training data of different training surfaces. The "Ca", "Mg", and "Mica" labels represent the training sets of (10$\bar{1}$4) calcite, (10$\bar{1}$4) magnesite, and (001) mica surfaces, respectively. The predictions are shown on the calcite test set, the magnesite test set, the mica test set, and the surfaces of dolomite, pyrophyllite, and boehmite.

the MAE of the DL models, we can infer that the energy values above ~5.1 indicate a reasonable prediction from our DL U-Net model. This implies that DL prediction can be used for the surfaces with an energy score above ~5.1, otherwise MD simulation should be performed, cif Fig. 1 (c).

4. Conclusion

The DL U-Net architecture is a robust network suitable for machine translation tasks. The U-Net, combined with the novel four-channel forcefield-based descriptor — introduced in this work — can predict the hydration layers over surfaces with varied atomic structure. This element-agnostic approach is rapid, robust, and reliable. Additionally, the transfer learning technique is demonstrated during training such network. With the ability to reuse the training sets and the pre-trained model, the process of augmenting the training surfaces with newer surfaces is made straightforward. The energy based DL network is shown to identify out-of-distribution surfaces with its score as a figure-of-merit. A possible workflow is identified where new surfaces are tested with the energy based DL network, and the DL network is used to predict the hydration layers over the surface if the score is above the threshold, otherwise an MD simulation is performed. Moreover, this MD simulation of the out-of-distribution surface can be included in the training data to make the DL network more robust. This makes the workflow highly scalable and applicable to other surfaces in general.

Further work can involve using the solvent tip approximation (STA)[19], or its successor extended-STA[21,17], to simulate hydration layers as seen by AFM. A network trained on such target hydration layers can circumvent an additional step of performing such simulations, and get closer to direct AFM characterisation with a comprehensive DL workflow. Ultimately, as new materials are introduced to the training workflow, the DL network should gradually develop into a general tool for the rapid prediction of hydration structures on any interface without the need for direct MD simulations.

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

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at https://doi.org/10.1016/j.molliq.2022.120571.

References

[1] K. Wandelt, S. Thurgate, (Eds.) Solid-Liquid Interfaces: Macroscopic Phenomena — Microscopic Understanding. Topics in Applied Physics (Springer-Verlag, 2003). https://www.springer.com/gp/book/9783540425830.

[2] F. Zaera, Probing liquid/solid interfaces at the molecular level, Chem. Rev. 112 (2012) 2920–2986, URL https://pubs.acs.org/doi/10.1021/cr2002068.

[3] G.E. Brown et al., Metal oxide surfaces and their interactions with aqueous solutions and microbial organisms, Chem. Rev. 99 (1999) 77–174.

[4] Binnig, G., Gerber, C., Stoll, E., Albrecht, T. & Quate, C. Atomic resolution with atomic force microscope. Europhysics Letters 3, 1281 (1987). Publisher: IOP Publishing.

[5] T. Fukuma, T. Ichii, K. Kobayashi, H. Yamada, K. Matsushige, True-molecular resolution imaging by frequency modulation atomic force microscopy in various environments, Appl. Phys. Lett. 86 (2005) 034103.

[6] T. Fukuma, Y. Ueda, S. Yoshioka, H. Asakawa, Atomic-scale distribution of water molecules at the mica-water interface visualized by three-dimensional scanning force microscopy, Phys. Rev. Lett. 104 (2010) 016101, URL https://link.aps.org/doi/10.1103/PhysRevLett.104.016101.

[7] K. Voïtchovsky et al., Direct mapping of the solid-liquid adhesion energy with subnanometre resolution, Nat. Nanotechnol. 5 (2010) 401.

[8] S. Rode, N. Oyabu, K. Kobayashi, H. Yamada, A. Kuehnle, True Atomic-Resolution Imaging of (10(1)over-bar4) Calcite in Aqueous Solution by Frequency Modulation Atomic Force Microscopy, Langmuir 25 (2009) 2850–2853.

[9] I. Siretanu, D. van den Ende, F. Mugele, Atomic structure and surface defects at mineral-water interfaces probed by in situ atomic force microscopy, Nanoscale 8 (2016) 8220–8227.

[10] K. Umeda et al., Atomic-resolution three-dimensional hydration structures on a heterogeneously charged surface, Nat. Commun. 8 (2017) 2111.

[11] Y.F. Dufrene et al., Imaging modes of atomic force microscopy for application in molecular and cell biology, Nat. Nanotechnol. 12 (2017) 295–307.

[12] T. Fukuma, R. Garcia, Atomic- and Molecular-Resolution Mapping of Solid-Liquid Interfaces by 3d Atomic Force Microscopy, ACS Nano 12 (2018) 11785–11797, https://doi.org/10.1021/acsnano.8b07216.

[13] Y. Teduka, A. Sasahara, H. Onishi, Atomic Force Microscopy Imaging of Crystalline Sucrose in Alcohols, ACS Omega 5 (2020) 2569–2574, https://doi.org/10.1021/acsomega.9b02660.

[14] H. Söngen et al., Resolving Point Defects in the Hydration Structure of Calcite (10.4) with Three-Dimensional Atomic Force Microscopy, Phys. Rev. Lett. 120 (2018) 116101, URL https://link.aps.org/doi/10.1103/PhysRevLett.120.116101.

[15] M. Watkins, A.L. Shluger, Mechanism of contrast formation in atomic force microscopy in water, Phys. Rev. Lett. 105 (2010) 196101.

[16] T. Fukuma et al., Mechanism of atomic force microscopy imaging of three-dimensional hydration structures at a solid-liquid interface, Physical Review B 92 (2015) 155412, URL https://link.aps.org/doi/10.1103/PhysRevB.92.155412.

[17] E. Nakouzi et al., Moving beyond the Solvent-Tip Approximation to Determine Site-Specific Variations of Interfacial Water Structure through 3D Force Microscopy, The Journal of Physical Chemistry C 125 (2021) 1282–1291, URL https://pubs.acs.org/doi/10.1021/acs.jpcc.0c07901.

[18] M.R. Uhlig, S. Benaglia, R. Thakkar, J. Comer, R. Garcia, Atomically resolved interfacial water structures on crystalline hydrophilic and hydrophobic surfaces, Nanoscale 13 (2021) 5275–5283, URL http://xlink.rsc.org/?DOI=D1NR00351H.

[19] M. Watkins, B. Reischl, A simple approximation for forces exerted on an AFM tip in liquid, J. Chem. Phys. 138 (2013) 154703.

[20] K.-I. Amano, K. Suzuki, T. Fukuma, O. Takahashi, H. Onishi, The relationship between local liquid density and force applied on a tip of atomic force microscope: A theoretical analysis for simple liquids, Journal Of Chemical Physics 139 (2013) 224710.

[21] K. Miyazawa et al., Tip dependence of three-dimensional scanning force microscopy images of calcite–water interfaces investigated by simulation and experiments, Nanoscale 12 (2020) 12856–12868, URL https://pubs.rsc.org/en/content/articlelanding/2020/nr/d0nr02043e.

[22] L. Gross, F. Mohn, N. Moll, P. Liljeroth, G. Meyer, The Chemical Structure of a Molecule Resolved by Atomic Force Microscopy, Science 325 (2009) 1110–1114.

[23] N. Pavliček, L. Gross, Generation, manipulation and characterization of molecules by atomic force microscopy, Nature Reviews Chemistry 1 (2017) 1–11, URL https://www.nature.com/articles/s41570-016-0005.

[24] T. Ando, High-Speed Atomic Force Microscopy in Biology: Directly Watching Dynamics of Biomolecules in Action, NanoScience and Technology, Springer, Berlin Heidelberg, 2022, URL https://link.springer.com/10.1007/978-3-662-64785-1.

[25] Y.S. Ranawat, Y.M. Jaques, A.S. Foster, Predicting hydration layers on surfaces using deep learning, Nanoscale Advances 3 (2021) 3447–3453, URL http://xlink.rsc.org/?DOI=D1NA00253H.

[26] Zhuang, F. et al. A comprehensive survey on transfer learning. Proceedings of the IEEE 109, 43–76 (2020). Publisher: IEEE.

[27] Momma, K. & Izumi, F. VESTA 3 for three-dimensional visualization of crystal, volumetric and morphology data. Journal of applied crystallography 44, 1272–1276 (2011). Publisher: International Union of Crystallography.

[28] J. Schmidhuber, Deep learning in neural networks: An overview, Neural Networks 61 (2015) 85–117, URL https://linkinghub.elsevier.com/retrieve/pii/S0893608014002135.

[29] Grathwohl, W. et al. Your Classifier is Secretly an Energy Based Model and You Should Treat it Like One. arXiv preprint arXiv:1912.03263 (2020). http://arxiv.org/abs/1912.03263.

[30] S.V. Kalinin, B.G. Sumpter, R.K. Archibald, Big–deep–smart data in imaging for guiding materials design, Nat. Mater. 14 (2015) 973–980.

[31] S.V. Kalinin et al., Big, Deep, and Smart Data in Scanning Probe Microscopy, ACS Nano 10 (2016) 9068–9086.

[32] M. Ziatdinov, A. Maksov, S.V. Kalinin, Learning surface molecular structures via machine vision, npj Comput. Mater. 3 (2017) 31.

[33] O.M. Gordon, P.J. Moriarty, Machine learning at the (sub)atomic scale: Next generation scanning probe microscopy, Machine Learning: Science and Technology 1 (2020) 023001, URL https://doi.org/10.1088%2F2632-2153%2Fab7d2f.

[34] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, Nature 521 (2015) 436–444.

[35] I. Azuri, I. Rosenhek-Goldian, N. Regev-Rudzki, G. Fantner, S.R. Cohen, The role of convolutional neural networks in scanning probe microscopy: a review, Beilstein Journal of Nanotechnology 12 (2021) 878–901, URL https://www.beilstein-journals.org/bjnano/articles/12/66.

[36] B. Alldritt et al., Automated structure discovery in atomic force microscopy, Science Advances 6 (2020) eaay6913, URL https://advances.sciencemag.org/content/6/9/eaay6913.

[37] O.M. Gordon, J.E.A. Hodgkinson, S.M. Farley, E.L. Hunsicker, P.J. Moriarty, Automated Searching and Identification of Self-Organized Nanostructures, Nano Lett. 20 (2020) 7688–7693, https://doi.org/10.1021/acs.nanolett.0c03213.

[38] J. Croshaw, T. Dienel, T. Huff, R. Wolkow, Atomic defect classification of the H-Si(100) surface through multi-mode scanning probe microscopy, Beilstein Journal of Nanotechnology 11 (2020) 1346–1360, URL https://www.beilstein-journals.org/bjnano/articles/11/119.

[39] M. Rashidi et al., Deep learning-guided surface characterization for autonomous hydrogen lithography, Machine Learning: Science and Technology 1 (2020) 025001, URL https://iopscience.iop.org/article/10.1088/2632-2153/ab6d5e.

[40] J.C. Thomas et al., Autonomous scanning probe microscopy investigations over WS2 and au111, npj Computational Materials 8 (2022) 99, URL https://www.nature.com/articles/s41524-022-00777-9.

[41] N. Oinonen et al., Electrostatic discovery atomic force microscopy, ACS Nano 16 (2022) 89–97, URL https://pubs.acs.org/doi/10.1021/acsnano.1c06840.

[42] M. Rashidi, R.A. Wolkow, Autonomous Scanning Probe Microscopy in SituTip Conditioning through Machine Learning, ACS nano 12 (2018) 5185–5189.

[43] B. Alldritt et al., Automated tip functionalization via machine learning in scanning probe microscopy, Comput. Phys. Commun. 273 (2022) 108258, URL https://linkinghub.elsevier.com/retrieve/pii/S0010465521003702.

[44] S.V. Kalinin et al., Automated and autonomous experiments in electron and scanning probe microscopy, ACS Nano 15 (2021) 12604–12627, URL https://pubs.acs.org/doi/10.1021/acsnano.1c02104.

[45] A. Krull, P. Hirsch, C. Rother, A. Schiffrin, C. Krull, Artificial-intelligence-driven scanning probe microscopy, Communications Physics 3 (2020) 54, URL http://www.nature.com/articles/s42005-020-0317-3.

[46] J. Sotres, H. Boyd, J.F. Gonzalez-Martinez, Enabling autonomous scanning probe microscopy imaging of single molecules with deep learning, Nanoscale 13 (2021) 9193–9203, URL http://xlink.rsc.org/?DOI=D1NR01109J.

[47] K. Kim et al., Molecular features of hydration layers: Insights from simulation, microscopy, and spectroscopy, The Journal of Physical Chemistry C 126 (2022) 8967–8977, URL https://pubs.acs.org/doi/10.1021/acs.jpcc.2c01313.

[48] P. Raiteri, R. Demichelis, J.D. Gale, Thermodynamically consistent force field for molecular dynamics simulations of alkaline-earth carbonates and their aqueous speciation, The Journal of Physical Chemistry C 119 (2015) 24447–24458, URL http://pubs.acs.org/doi/10.1021/acs.jpcc.5b07532.

[49] B. Reischl, P. Raiteri, J.D. Gale, A.L. Rohl, Can point defects in surfaces in solution be atomically resolved by atomic force microscopy?, Phys Rev. Lett. 117 (2016) 226101, URL https://link.aps.org/doi/10.1103/PhysRevLett.117.226101.

[50] Y. Wu, H.L. Tepper, G.A. Voth, Flexible simple point-charge water model with improved liquid-state properties, J. Chem. Phys. 124 (2006) 024503, URL http://aip.scitation.org/doi/10.1063/1.2136877.

[51] R.T. Cygan, J.-J. Liang, A.G. Kalinichev, Molecular models of hydroxide, oxyhydroxide, and clay phases and the development of a general force field, J. Phys. Chem. B 108 (2004) 1255–1266, URL https://pubs.acs.org/doi/10.1021/jp0363287.

[52] W.L. Jorgensen, J. Chandrasekhar, J.D. Madura, R.W. Impey, M.L. Klein, Comparison of simple potential functions for simulating liquid water, J. Chem. Phys. 79 (1983) 926–935, URL http://aip.scitation.org/doi/10.1063/1.445869.

[53] S.R. van Lin, K.K. Grotz, I. Siretanu, N. Schwierz, F. Mugele, Ion-specific and pH-dependent hydration of mica–electrolyte interfaces, Langmuir 35 (2019) 5737–5745, URL http://pubs.acs.org/doi/10.1021/acs.langmuir.9b00520.

[54] S. Plimpton, Fast parallel algorithms for short-range molecular dynamics, Journal of computational physics 117 (1995) 1–19, URL https://lammps.sandia.gov/index.html.

[55] I. Goodfellow, Y. Bengio, A. Courville, Deep Learning, MIT Press, 2016.

[56] Ronneberger, O., Fischer, P. & Brox, T. U-net: Convolutional networks for biomedical image segmentation. In International Conference on Medical image computing and computer-assisted intervention, 234–241 (Springer, 2015).

[57] Song, Y. & Ermon, S. Generative Modeling by Estimating Gradients of the Data Distribution. arXiv:1907.05600 [cs, stat] (2020). http://arxiv.org/abs/1907.05600. ArXiv: 1907.05600.

[58] Carreira-Perpinan, M.A. & Hinton, G. On contrastive divergence learning. In International workshop on artificial intelligence and statistics, 33–40 (PMLR, 2005).

[59] J.D. Hunter, Matplotlib: A 2d graphics environment, Computing In Science & Engineering 9 (2007) 90–95.

[60] H. Söngen et al., Chemical identification at the solid–liquid interface. Langmuir 33 (2017) 125–129. Publisher: ACS Publications.