Electrostatic Discovery Atomic Force Microscopy

Niko Oinonen,[1*], Chen Xu,[1*], Benjamin Alldritt,[1*]
Filippo Federici Canova,[1,2] Fedor Urtev,[1,3] Shuning Cai,[1] Ondřej Krejčí,[1]
Juho Kannala,[3] Peter Liljeroth[1†] and Adam S. Foster[1,4†]

[1]Department of Applied Physics, Aalto University, 00076 Aalto, Helsinki, Finland

[2]Nanolayers Research Computing Ltd, London N12 0HL, United Kingdom

[3]Department of Computer Science, Aalto University, 00076 Aalto, Helsinki, Finland

[4]WPI Nano Life Science Institute (WPI-NanoLSI), Kanazawa University,
Kakuma-machi, Kanazawa 920-1192, Japan

*These authors contributed equally.

†To whom correspondence should be addressed;
E-mail: peter.liljeroth@aalto.fi; adam.foster@aalto.fi.

Abstract

While offering high resolution of atomic and electronic structure, Scanning Probe Microscopy techniques have found greater challenges in providing reliable electrostatic characterization at the same scale. In this work, we offer Electrostatic Discovery Atomic Force Microscopy, a machine learning based method which provides immediate maps of the electrostatic potential directly from Atomic Force Microscopy images with functionalized tips. We apply this to characterize the electrostatic properties of a variety of molecular systems and compare directly to reference simulations, demonstrating good agreement. This approach offers reliable atomic scale electrostatic maps on any system with minimal computational overhead.

Keywords: atomic force microscopy, machine learning, tip functionalization, chemical identification, electrostatics

Introduction

The electrostatic properties of molecules are dominant in a wide variety of processes and technologies, from catalysis and chemical reactions, [1] to molecular electronics [2] and biological functions. [3] In general, if we can understand the link between molecular function and electrostatics, it offers powerful tools to control and design functionality with nanoscale precision. [4] At this scale, Scanning Probe Microscopy (SPM) is the characterization technique of choice and Scanning Tunneling Microscopy (STM) has become the engine of local electronic characterization for conducting systems, [5, 6] while AFM is a general tool for nanoscale imaging without material restrictions. [7, 8] In high-resolution studies, AFM has evolved from its origins [9] into a breakthrough technique in studies of molecular systems. [10, 11] This has been driven by the use of functionalized tips, and AFM in ultra high vacuum (UHV) now offers a window into molecular structure on surfaces – aside from the detailed resolution of the results of molecular assembly, it is possible to study bond order, charge distributions and the individual steps of on-surface chemical reactions. [11] More recently, additions to the SPM family such as alternate-charging STM [12] and single-electron transfer AFM [13] have offered approaches to study charge behaviour in molecular systems.

While all these methods give indirect information on the electrostatic properties of the system being studied, significant efforts have been made to develop systematic techniques to directly characterize electrostatic properties. In particular, Kelvin Probe Microscopy (KPFM) was introduced [14, 15] to simultaneously explore the topography and local contact potential difference with atomic resolution. Despite success in characterizing the electrostatic properties of surfaces, [16–19] and even proteins, [20] the technique has limitations that prevent widespread adoption. Generally, it is experimentally challenging, requiring much longer measurement times than equivalent STM or AFM experiments and can be prone to tip-convolutions.

[21] More generally, the varying contributions to the signal mean it is very challenging to obtain quantitative measurements from KPFM. [22] The step-change in molecular characterization offered by functionalized tips in AFM has also been harnessed for electrostatic analysis, with KPFM being applied with functionalized tips to provide a local potential maps of single molecules. [23, 24] However, the outstanding challenges of KPFM remain, or are even exaggerated: there is no rigorous KPFM theory at the atomic scale, the usually assumed qualitative proportionality to the out-of-plane electric field gradient breaks down at small tip-sample separations and convolution with unknown background force contributions further complicates the analysis. Attempts to address these problems led to the recent development of scanning quantum dot microscopy, [25] which offers an alternative approach to map the local potential of a surface and adsorbates in high resolution. [26] While powerful, the technique relies on quantum dot tip functionalization and a dedicated controller, again limiting its broad implementation as yet.

In this work, we were inspired by earlier efforts which use AFM tips functionalized with molecules of different electrostatic character to resolve the electrostatic potential by comparing the characteristic distortions of each tip. [27] However, the limitations of the approach meant it was unable to provide quantitative accuracy. Here, we offer Electrostatic Discovery Atomic Force Microscopy (ED-AFM), a machine learning (ML) approach that can predict accurate electrostatic fields directly from a set of standard experimental AFM images. This methodology offers convenient access to the electrostatic potential of molecules adsorbed on surfaces, which will be important in - for example - understanding their catalytic activity, identifying products of on-surface synthesis routes and facilitating chemical identification of functional groups in unknown molecules.

Results and Discussion

At the heart of the ED-AFM methodology lies a convolutional neural network that is trained to connect input data – a set of constant height AFM images of the frequency shift Δf at different tip-sample distances acquired with two different tip charges – to a map of the electrostatic potential over the molecule (ES map descriptor). The details of the procedure are given in the Methods section.

Benchmark systems

In order to benchmark the ED-AFM method, we first consider three molecular systems using only simulated data. The first two, "N2-(2-Chloroethyl)-N-(2,6-dimethylphenyl)-N2-methyl-glycinamide (NCM)" and "4-Pyridinecarboxylic acid, 2-[(1E)-2-thienylmethylene]hydrazide (PTH)" (see Figs. 1A,B) were chosen due to the presence of different functional groups and bonds, as well as due to their non-planar geometry. In the third example, we focus on a charge transfer complex, polar tetrathiafulvalene thiadiazole (TTF-TDZ) (see Fig. 1C), which was characterized electrostatically in a previous work using KPFM and *ab initio* simulations. [24] These examples are only considered as free-standing molecules, and therefore the presented orientations and geometries are likely not fully representative of those that the molecules would adopt on a surface.

We consider here the predictions in comparison to the point-charge reference that the model was trained to reproduce. In all cases the match between the predicted and the reference descriptors is generally excellent. The positive and negative regions are predicted in the correct places at correct magnitude with some small imprecision at the edge regions. For a more quantitative comparison, we consider a relative error metric

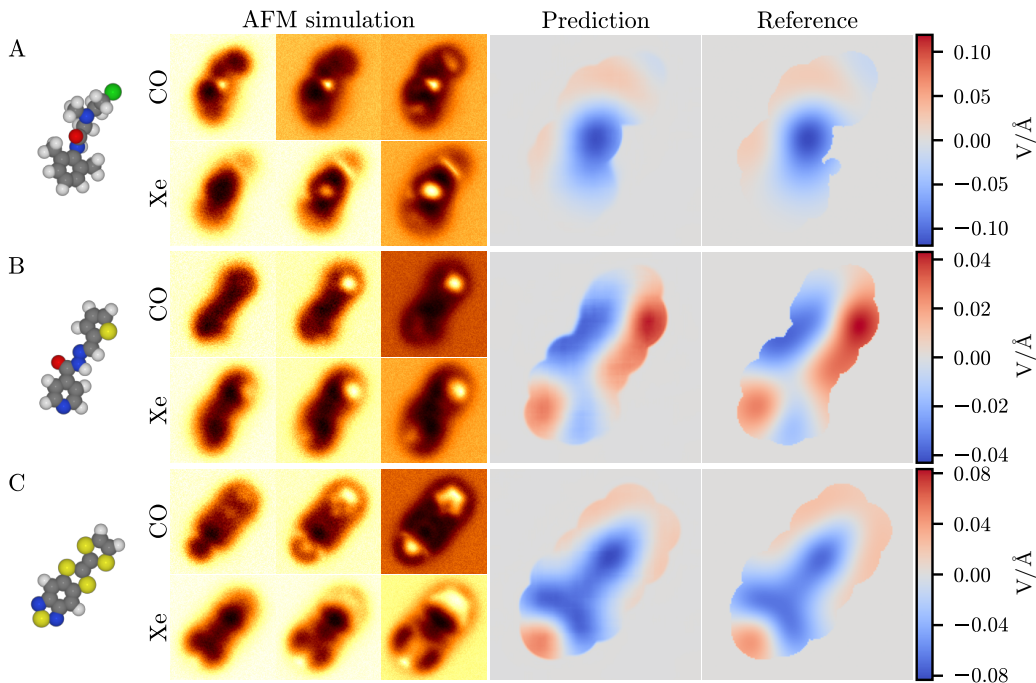$$\frac{\sum_i |y_i - \tilde{y}_i| / N}{\max \tilde{y} - \min \tilde{y}},$$

Figure 1: Predictions on simulated AFM images. Predictions are shown for three test systems, (**A**) N2-(2-Chloroethyl)-N-(2,6-dimethylphenyl)-N2-methylglycinamide, (**B**) 2-[(1E)-2-thienylmethylene]-hydrazide, and (**C**) tetrathiafulvalene thiadiazole. In each case are shown, from left to right, the 3D structure of the molecule, three out of six input AFM images at different tip-sample distances for both tip functionalizations, and the predicted and reference ES Map descriptors. The colorbar scale for the prediction and the reference is the same on each row.

where y is the predicted descriptor, $\tilde{y}$ is the reference descriptor, and the sum is over the N pixels in the descriptor. This is the mean absolute error in the prediction relative to the range of values in the reference. For our three benchmark examples, we find that the relative errors are 1.04% for NCM, 1.79% for PTH, and 2.29% for TTF-TDZ.

Validation

Clearly, the real test of ED-AFM is with experimental data and we consider three representative example molecular systems. In our first example, we consider perylenetetracarboxylic dian-hydride (PTCDA) on the Cu(111) surface (see Fig. 2), a benchmark system in the analysis of

5

Figure 2: Comparison of simulated and experimental predictions for perylenetetracarboxylic dianhydride. On the left are shown three out of six input AFM images at different tip-sample distances for both tip functionalizations, and on the right are the model predictions for both simulation and experiment and the reference descriptor. Both predictions and the reference are on the same colorbar scale. The molecule geometry used in the simulation is shown on the bottom right.

molecules on metal surfaces and in characterizing electrostatic interface properties. [26, 28, 29] The simulation is done on a free-standing molecule with planar geometry using point-charges for electrostatics. In the reference ES Map descriptor we find that the ends of the molecule with the three oxygens has a negative field, as would be expected by the electronegativity of oxygen, and the field in the middle of the molecule is positive. In line with the previous simulation examples, the prediction based on the simulated AFM images is in good agreement with the reference. The prediction from the experimental AFM images similarly shows negative field over the ends of the molecule and positive field in between. This matches well with the reference, except for the somewhat weaker magnitude of the field in the experimental prediction.

We note that in the experimental Xe-AFM images, there appears an abnormally large bright feature in the upper right corner over one of the oxygens. The origin of this artifact is unclear, but we speculate that it could be due to a difference in charge at that site (see SI Sec. "Possible extra electron in PTCDA" for further discussion). Despite this unusual feature in the AFM

images, the model prediction does not appear to be greatly disturbed over the corresponding region. Another unusual feature in this set of images is that there is a gradient in the background of the AFM images decreasing from the upper edge towards the lower edge of the images. This is due to the experiment being performed at a slight tilt with respect to the surface. Originally our model could not handle this feature in the images very well, since such features originating from the surface are normally absent from the simulations that only consider free molecules. After adding artificial gradients in random directions to the simulation images in training, we found that the model begun to perform more consistently on the PTCDA experiment. See SI Sec. "Surface tilt effect on model predictions" for further discussion.

For our second example, we study 1-bromo-3,5-dichlorobenzene (BCB) on the Cu(111) surface (see Fig. 3), a planar molecule with mixed halide functionalization. As with PTCDA, for simulations of this molecule we consider a completely planar free-standing molecule with point-charge electrostatics. In the reference ES Map descriptor we find a negative field both over the chlorines and in the middle of the molecule, close to neutral over the bromine, and positive over the hydrogens. Again, the simulation prediction is in good agreement with the reference. In the prediction for the experimental AFM images, we see a region of positive field running along the edge of the lower left part of the molecule, where we suppose the bromine is, and negative field over the other two halides. This matches quite well with the reference, also in the magnitude of the field, except for the missing positive region over the hydrogen opposing the bromine. We note here that the experimental Xe-AFM images have been flipped left-to-right and slightly rotated due to the molecule having rotated between the CO and Xe experiments. The original images can be seen in SI Fig. S6.

Our final example is a cluster of seven water molecules on the Cu(111) surface (see Fig. 4). Again, we use point-charge electrostatics, but this time we include the metallic surface, since this configuration of water molecules is only stable on a surface. The seven water molecules
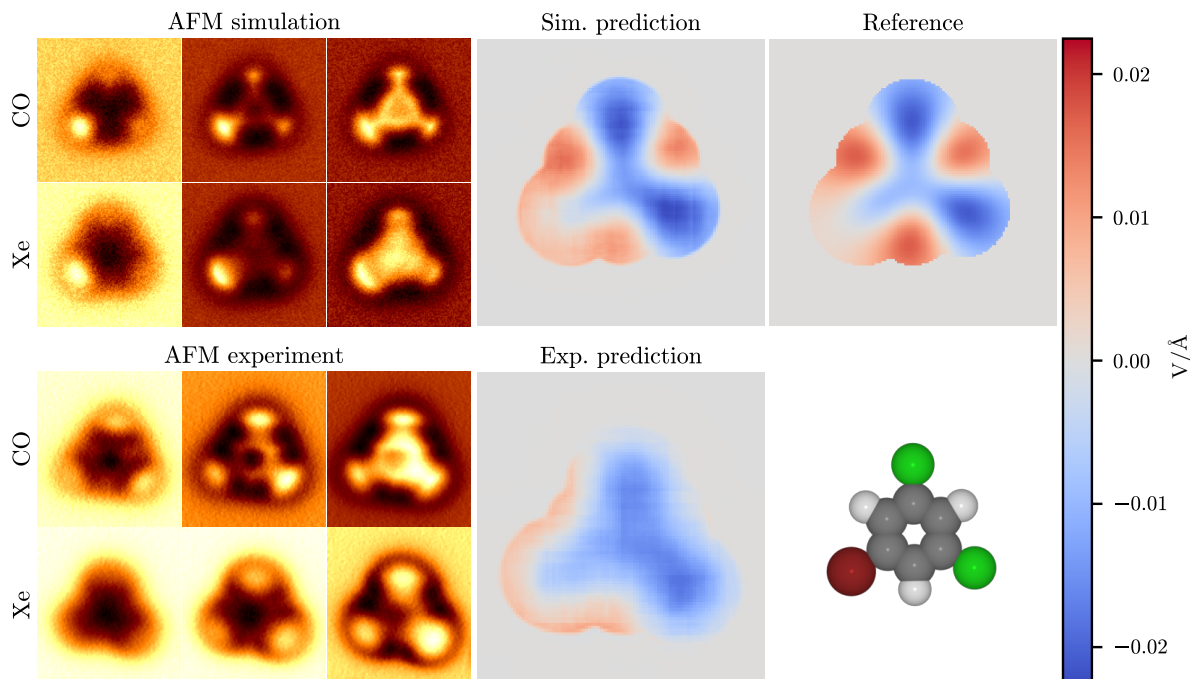
7

Figure 3: Comparison of simulated and experimental predictions for 1-bromo-3,5-dichlorobenzene. On the left are shown three out of six input AFM images at different tip-sample distances for both tip functionalizations, and on the right are the model predictions for both simulation and experiment and the reference descriptor. Both predictions and the reference are on the same colorbar scale. The molecule geometry used in the simulation is shown on the bottom right.

form a single five-member ring with additional two molecules forming an incomplete second ring. Similar water clusters with 9 and 10 molecules forming two complete rings were previously studied with a combination of DFT calculations and STM experiments. [30] In our calculations, we find that having the second ring be incomplete results in a better match between the simulated and experimental AFM images.

Here in the reference ES Map descriptor we find that the field is mostly positive over the whole structure with some neutral regions on the sides. The simulation prediction matches the reference well except for the sides where the prediction is more negative. The match of the experimental prediction with the reference is reasonable, capturing the positive regions at the top part and at the bottom left "leg" of the structure, but it also has significantly larger negative

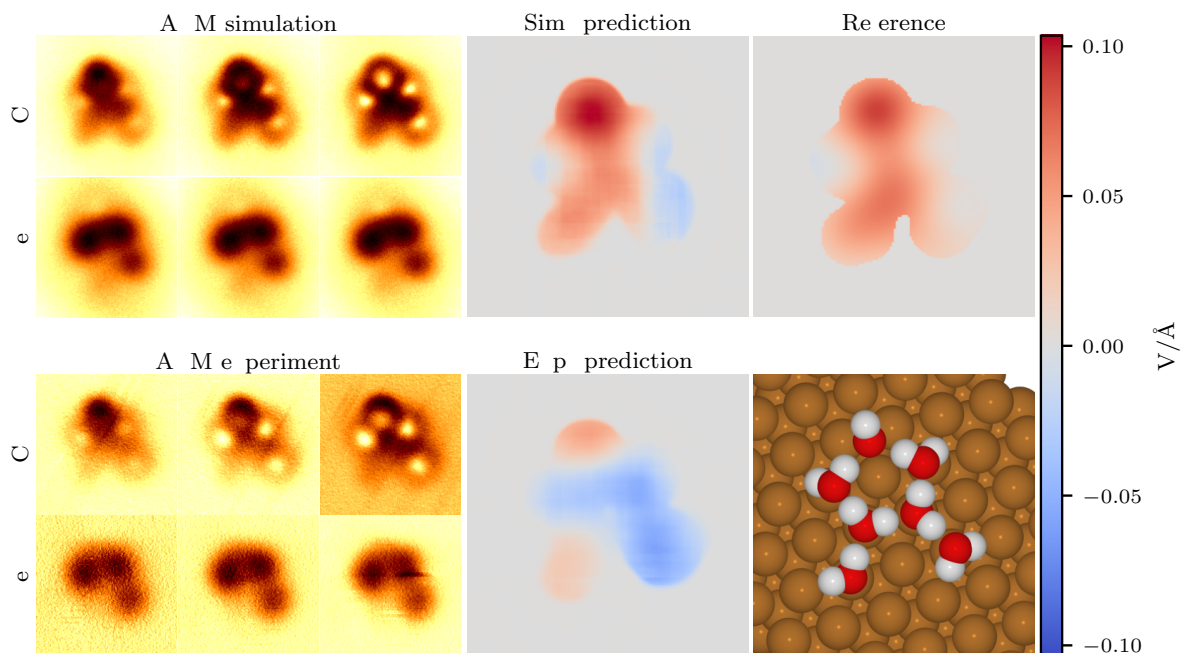region in the middle and right side compared to either the simulation prediction or the reference.

Figure 4: Comparison of simulated and experimental predictions for a water cluster on Cu(111). On the left are shown three out of six input AFM images at different tip-sample distances for both tip functionalizations, and on the right are the model predictions for both simulation and experiment and the reference descriptor. Both predictions and the reference are on the same colorbar scale. The molecule geometry used in the simulation is shown on the bottom right.

Since we have obtained more than the required 6 height-slices for each tip in the experiments, we also consider what happens to the predictions as functions of tip-sample distance for the PTCDA and BCB experiments and find that the predictions stay consistent over small deviations to the distance in either direction (see SI Sec. "Distance dependence").

Limitations of the current model

Until this point we have been considering only the point-charge electrostatics that the ML model was trained on, and found that the model performs well on the simulation examples and the experimental predictions are in fairly good agreement with the ones for simulations. However,

9

Figure 5: Prediction and reference for on-surface geometry of (**A**) BCB and (**B**) water cluster using the DFT Hartree potential for electrostatics in the AFM simulations and for the reference ES Map descriptor.

the point-charge model of electrostatics has its limitations and in many cases does not perfectly represent the true charge distribution and, by extension, the electric field of the sample. In order to test the validity of the results so far, we performed density functional theory (DFT) calculations using the all-electron density functional theory code FHI-AIMS, [31] implementing the "tight" basis with the PBE functional [32] and TS van der Waals [33] for all the test systems to obtain their Hartree potentials, which can be used for more accurate electrostatics in both the AFM simulations and the reference ES Map descriptors. We first test the Hartree potentials on our three simulation benchmark systems (Fig. S7 in SI) and find that the Hartree reference descriptors remain mostly the same as the point-charge ones. The predictions are still qualitatively fairly good, and are at least semi-quantitative in accuracy. The relative errors are 6.12%, 4.78%, and 7.34% for NCM, PTH, and TTF-TDZ, respectively.

Next, we test the Hartree potential on BCB, one of our experimental test systems. For this test we first relax the geometry on a Cu(111) surface to obtain a more accurate geometry for

the molecule. The resulting AFM simulation based on the Hartree potential rather than point charges, the ML prediction, and the reference descriptor are shown in Fig. 5A. In the on-surface geometry the bromine is attracted closer to the surface, giving the molecule a slight tilt. This makes the bromine appear less bright in the AFM simulation compared to the planar geometry with point charges, which matches better with the experimental AFM images. The Hartree reference ES Map has similar pattern of field as the point-charge reference over the edges of the molecule, but they differ in the middle over the carbon ring where the Hartree reference has strong positive field instead of negative field. The prediction on the simulation using the Hartree potential correctly catches this positive field in the middle of the molecule, but it misses the negative field over the chlorines and the magnitude of the field is overall too small, roughly by a factor of two. However, the prediction on the experimental AFM images compares more favourably with the ES Map based on point-charge than with DFT/Hartree, even though the Hartree potential from DFT should represent a more accurate description of reality. Note here that it is not *a priori* clear which (if either) ES map should the experiment reproduce: using point charges and the geometry of an isolated molecule is not expected to be the best description of the real BCB molecule on Cu(111) surface. On the other hand, if we run the AFM simulations based on DFT/Hartree molecular structure and electrostatics, we would not expect the ML to reproduce the ES map reference as the model is trained on point charges. Understanding this in detail requires development of a model trained on out-of-distribution samples using the Hartree potential, a focus of future work.

We also perform the DFT calculation for on-surface PTCDA, and find that in this case the model performs poorly even on the simulated data (Fig. S8 in SI). However, establishing accurate geometries for PTCDA on metal surfaces is very challenging, [34] and indeed, when we perform AFM simulations for our on-surface geometry for PTCDA (SI Fig. S8) we find that the simulation shows an asymmetric and much larger contrast between the ends and the middle

of the molecule than in the experimental images. This indicates that the real geometry of the molecule is more planar and more symmetric than in our DFT simulation. Therefore, we feel the reference ES Map descriptor for PTCDA obtained from the DFT Hartree potential here does not provide a valid point of comparison with the experimental prediction.

Finally, we test the Hartree potential on the water cluster (Fig. 5B) using the same geometry as with the point charges. Compared to BCB, the difference between the point charges and the Hartree potential is less pronounced, both in the simulated AFM images and the reference ES Map. The biggest difference is the appearance of a stronger positive region at the top of the structure and more negative field on the sides, which takes the reference, at least qualitatively, closer to the ES Map predicted from the experimental AFM images. The prediction on the simulated images using the Hartree potential is qualitatively good, but is weaker in magnitude, especially at the positive region at the top. The relatively weaker magnitude of the predictions compared to the reference seems to be the general pattern for all of our cases using the Hartree potential.

Conclusions

This work demonstrates that ED-AFM offers a method for the rapid prediction of electrostatic properties directly from experimental AFM images. We have shown that these predictions demonstrate quantitative accuracy with minimal computational cost once the machine learning infrastructure is trained. In particular, the comparison between reference electrostatics and ML predictions from simulated AFM data has an error of about 1-2%. However, it is also clear, as for any ML-based approach, that the method cannot predict what it has not learned and it performs poorly in systems where point charge electrostatics are a poor approximation - even for simulated data, the error increases to 5-7%. We will address this in future work: as part of the database generation, we have stored full density matrices with higher-order quantum

chemical accuracy [35, 36] and we are developing methods for their efficient incorporation into the training process. However, we also note that fixed point charges offer decent accuracy in many systems and are routinely used in molecular modelling. [37, 38] The further limitations of ED-AFM lie mainly in the challenges posed for experiments and, in particular, obtaining images of the same system with two different tips. As we have demonstrated, this is feasible on well-defined metal surfaces, but can pose problems on less standard samples. This can be at least partially alleviated by developing methods for the autonomous functionalization of the tip in AFM (see, *e.g.*, https://github.com/SINGROUP/Auto-CO-AFM). We note that the two different tips do not need to be CO and Xe, and we have shown, at least for simulated data, that other pairs would work effectively for ED-AFM. Finally, as for any machine learning based approach, the predictions can be confused by unexpected artefacts. We apply a suite of tools to improve robustness (see SI Secs. "Machine learning" and "AFM simulations") and due to the longer range of electrostatic forces, the predictions are less dominated by images at close approach and hence less sensitive to tip-induced relaxations. [39]

Finally we note that ED-AFM also offers the prospects of application beyond the examples considered here, to any system where electrostatic characterization is of interest. For example, expansion to predictions for assembled layers, [40, 41] defects [42, 43] and charge dynamics, [44] requires only development of the training data to ensure that it is general enough. The method can also be applied to images of systems where direct simulation is impossible due to size, complexity or lack of information. Combined with developments in the autonomous functionalization of the tip in SPM, [45] this promises a future of potential in electrostatics for ED-AFM.

Methods/Experimental

In general, the adoption of ML methods into materials analysis has seen rapid recent growth [46–48] and this has been followed by an equivalent growth in its applications to image analysis in SPM. [45, 49–56] Here we build upon our ML method for predicting molecular structure from AFM images, [39] to predict the electrostatic field of the sample molecule.


Figure 6: Schematic of the ED-AFM method. We train a neural network that takes two sets of AFM images as input and translates them to the ES Map descriptor, which is the vertical component of the electrostatic field over the sample molecule. The model is trained on simulated sets of input-output pairs calculated from a database of several tens of thousands of molecule geometries. The trained model can then be applied to experimental AFM images to produce a prediction of the sample electric field.

The overall idea of our method is illustrated in Fig. 6. We train a deep learning model on a dataset of simulated AFM images and reference descriptors based on a large database of molecular geometries, including electrostatics at the level of point charges taken from the associated quantum chemistry calculations. [39] For more details on how the dataset is generated from the geometries, see Sec. "Data set" in the Supplementary Information (SI). The trained ML model can then be used to make predictions with experimental data as input. More specifically, the ML model takes as input two sets of six AFM images of the same sample at different tip-sample distances, imaged with two different functionalizations of the AFM tip. The model, which is an

14

Attention-U-Net-type convolutional neural network [57, 58] (more details in SI Sec. "Machine learning"), translates the AFM images into a descriptor of the imaged sample, which we call the Electrostatic (ES) Map. The ES Map is defined as the vertical component of the electrostatic field calculated on a constant-height surface $4\,\text{Å}$ above the highest atom in the sample. Furthermore, the non-zero region is cut only into the region where the sample molecule of interest is visible in the AFM images, such that the ML model is not asked to predict the field over the background where there are no discernible features present. For more detailed description of how the ES Map descriptor is generated, see SI Sec. "ES map descriptor".

The use of two sets of AFM images in the input is motivated by the observation that the different distortions in the images obtained with different tip functionalizations are linked to the different electronic charges on the tips. [27] Thus, given a database of such pairs of images, an ML model should be able to learn what role the electrostatics play in the formation of the images and separate the electrostatic contribution from other forces that contribute to the images. Here for the tip functionalizations we use CO, which has a slightly negative charge, and Xe, which has somewhat positive charge. Other choices for functionalization are possible, and we investigate the alternative tip combinations of Cl-CO and Cl-Xe on simulated data in SI Sec. "Other tip combinations". We also tried training a model using images of only a single tip functionalization with CO, but found the results to be less robust (see SI Sec. "Single-channel measurements"). As in our previous work, [39] we train the model using simulated AFM images and validate the model using both simulated and experimental AFM images. Our implementation of the model in Pytorch [59] with pretrained weights can be found at `https://github.com/SINGROUP/ED-AFM`.

Supporting Information

The Supporting Information contains details of the machine learning models and training data, including investigation of the sensitivity of the predictions to tip-surface distance, noise, various simulation parameters and the number of channels used in experiments. It also contains expanded discussion of the AFM simulations and experiments themselves.

Acknowledgements

References

(1) Ciampi, S.; Darwish, N.; M. Aitken, H.; Díez-Pérez, I.; L. Coote, M. Harnessing Electrostatic Catalysis in Single Molecule, Electrochemical and Chemical Systems: A Rapidly Growing Experimental Tool Box. *Chemical Society Reviews* **2018**, *47*, 5146–5164.

(2) Xiang, D.; Wang, X.; Jia, C.; Lee, T.; Guo, X. Molecular-Scale Electronics: From Concept to Function. *Chemical Reviews* **2016**, *116*, 4318–4440.

(3) Davis, M. E.; McCammon, J. A. Electrostatics in Biomolecular Structure and Dynamics. *Chemical Reviews* **1990**, *90*, 509–521.

(4) Beck, M. E.; Hersam, M. C. Emerging Opportunities for Electrostatic Control in Atomically Thin Devices. *ACS Nano* **2020**, *14*, 6498–6518.

(5)   Zandvliet, H. J.; van Houselt, A. Scanning Tunneling Spectroscopy. *Annual Review of Analytical Chemistry* **2009**, *2*, 37–55.

(6)   Ervasti, M. M.; Schulz, F.; Liljeroth, P.; Harju, A. Single- And Many-Particle Description of Scanning Tunneling Spectroscopy. *Journal of Electron Spectroscopy and Related Phenomena* **2017**, *219*, 63–71.

(7)   Loos, J. The Art of SPM: Scanning Probe Microscopy in Materials Science. *Advanced Materials* **2005**, *17*, 1821–1833.

(8)   Dufrene, Y. F.; Ando, T.; Garcia, R.; Alsteens, D.; Martínez-Martín, D.; Engel, A.; Gerber, C.; Müller, D. J. Imaging Modes of Atomic Force Microscopy for Application in Molecular and Cell Biology. *Nature Nanotechnology* **2017**, *12*, 295–307.

(9)   Giessibl, F. J. Advances in Atomic Force Microscopy. *Reviews Of Modern Physics* **2003**, *75*, 949–983.

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

(11) Pavliček, N.; Gross, L. Generation, Manipulation and Characterization of Molecules by Atomic Force Microscopy. *Nature Reviews Chemistry* **2017**, *1*, 0005.

(12) Patera, L. L.; Queck, F.; Scheuerer, P.; Repp, J. Mapping Orbital Changes upon Electron Transfer with Tunnelling Microscopy on Insulators. *Nature* **2019**, *566*, 245–248.

(13) Fatayer, S.; Albrecht, F.; Zhang, Y.; Urbonas, D.; Peña, D.; Moll, N.; Gross, L. Molecular Structure Elucidation with Charge-State Control. *Science* **2019**, *365*, 142–145.

(14) Nonnenmacher, M.; O'Boyle, M. P.; Wickramasinghe, H. K. Kelvin Probe Force Microscopy. *Applied Physics Letters* **1991**, *58*, 2921–2923.

(15) Sadewasser, S.; Glatzel, T.; Sadewasser, S.; Glatzel, T., *Kelvin Probe Force Microscopy: From Single Charge Detection to Device Characterization*; Springer: Cham, 2018; Vol. 65.

(16) Palermo, V.; Palma, M.; Samorì, P. Electronic Characterization of Organic Thin Films by Kelvin Probe Force Microscopy. *Advanced Materials* **2006**, *18*, 145–164.

(17) Sadewasser, S.; Jelinek, P.; Fang, C.-K.; Custance, O.; Yamada, Y.; Sugimoto, Y.; Abe, M.; Morita, S. New Insights on Atomic-Resolution Frequency-Modulation Kelvin-Probe Force-Microscopy Imaging of Semiconductors. *Physical Review Letters* **2009**, *103*, 266103.

(18) Gross, L.; Schuler, B.; Mohn, F.; Moll, N.; Pavliček, N.; Steurer, W.; Scivetti, I.; Kotsis, K.; Persson, M.; Meyer, G. Investigating Atomic Contrast in Atomic Force Microscopy and Kelvin Probe Force Microscopy on Ionic Systems Using Functionalized Tips. *Physical Review B* **2014**, *90*, 155455.

(19) Schulz, F.; Ritala, J.; Krejčí, O.; Seitsonen, A. P.; Foster, A. S.; Liljeroth, P. Elemental Identification by Combining Atomic Force Microscopy and Kelvin Probe Force Microscopy. *ACS Nano* **2018**, *12*, 5274–5283.

(20) Sinensky, A.; Belcher, A. Label-Free and High-Resolution Protein/DNA Nanoarray Analysis Using Kelvin Probe Force Microscopy. *Nature Nanotechnology* **2007**, *2*, 653–659.

(21) Zerweck, U.; Loppacher, C.; Otto, T.; Grafström, S.; Eng, L. M. Accuracy and Resolution Limits of Kelvin Probe Force Microscopy. *Physical Review B* **2005**, *71*, 125424.

(22) Söngen, H.; Rahe, P.; Neff, J. L.; Bechstein, R.; Ritala, J.; Foster, A. S.; Kühnle, A. The Weight Function for Charges—A Rigorous Theoretical Concept for Kelvin Probe Force Microscopy. *Journal Of Applied Physics* **2016**, *119*, 025304.

(23) Mohn, F.; Gross, L.; Moll, N.; Meyer, G. Imaging the Charge Distribution within a Single Molecule. *Nature Nanotechnology* **2012**, *7*, 227–231.

(24) Schuler, B.; Liu, S.-X.; Geng, Y.; Decurtins, S.; Meyer, G.; Gross, L. Contrast Formation in Kelvin Probe Force Microscopy of Single π-Conjugated Molecules. *Nano Letters* **2014**, *14*, 3342–3346.

(25) Wagner, C.; Green, M. F. B.; Leinen, P.; Deilmann, T.; Krüger, P.; Rohlfing, M.; Temirov, R.; Tautz, F. S. Scanning Quantum Dot Microscopy. *Physical Review Letters* **2015**, *115*, 026101.

(26) Wagner, C.; Green, M. F. B.; Maiworm, M.; Leinen, P.; Esat, T.; Ferri, N.; Friedrich, N.; Findeisen, R.; Tkatchenko, A.; Temirov, R.; Tautz, F. S. Quantitative Imaging of Electric Surface Potentials with Single-Atom Sensitivity. *Nature Materials* **2019**, *18*, 853–859.

(27) Hapala, P.; Švec, M.; Stetsovych, O.; van der Heijden, N. J.; Ondráček, M.; van der Lit, J.; Mutombo, P.; Swart, I.; Jelínek, P. Mapping the Electrostatic Force Field of Single Molecules from High-Resolution Scanning Probe Images. *Nature Communications* **2016**, *7*, 11560.

(28) Tautz, F. S. Structure and Bonding of Large Aromatic Molecules on Noble Metal Surfaces: The Example of PTCDA. *Progress in Surface Science* **2007**, *82*, 479–520.

(29) Burke, S. A.; LeDue, J. M.; Miyahara, Y.; Topple, J. M.; Fostner, S.; Grütter, P. Determination of the Local Contact Potential Difference of PTCDA on NaCl: A Comparison of Techniques. *Nanotechnology* **2009**, *20*, 264012.

(30) Liriano, M. L.; Gattinoni, C.; Lewis, E. A.; Murphy, C. J.; Sykes, E. C. H.; Michaelides, A. Water–Ice Analogues of Polycyclic Aromatic Hydrocarbons: Water Nanoclusters on Cu(111). *Journal of the American Chemical Society* **2017**, *139*, 6403–6410.

(31) Blum, V.; Gehrke, R.; Hanke, F.; Havu, P.; Havu, V.; Ren, X.; Reuter, K.; Scheffler, M. *Ab Initio* Molecular Simulations with Numeric Atom-Centered Orbitals. *Computer Physics Communications* **2009**, *180*, 2175–2196.

(32) Perdew, J. P.; Burke, K.; Ernzerhof, M. Generalized Gradient Approximation Made Simple. *Physical Review Letters* **1996**, *77*, 3865–3868.

(33) Tkatchenko, A.; Scheffler, M. Accurate Molecular van der Waals Interactions from Ground-State Electron Density and Free-Atom Reference Data. *Physical Review Letters* **2009**, *102*, 073005.

(34) Ruiz, V. G.; Liu, W.; Tkatchenko, A. Density-Functional Theory with Screened van der Waals Interactions Applied to Atomic and Molecular Adsorbates on Close-Packed and Non-Close-Packed Surfaces. *Physical Review B* **2016**, *93*, 035118.

(35) Crawford, T. D.; Schaefer III, H. F. In *Reviews in Computational Chemistry*; John Wiley & Sons, Ltd: Hoboken, 2007; Chapter 2, pp 33–136.

(36) Parrish, R. M.; Burns, L. A.; Smith, D. G. A.; Simmonett, A. C.; DePrince, A. E.; Hohenstein, E. G.; Bozkaya, U.; Sokolov, A. Y.; Di Remigio, R.; Richard, R. M.; Gonthier, J. F.; James, A. M.; McAlexander, H. R.; Kumar, A.; Saitow, M.; Wang, X.; Pritchard, B. P.; Verma, P.; Schaefer, H. F.; Patkowski, K., et al. Psi4 1.1: An Open-Source Electronic Structure Program Emphasizing Automation, Advanced Libraries, and Interoperability. *J. Chem. Theory Comput.* **2017**, *13*, 3185–3197.

(37) Cisneros, G. A.; Karttunen, M.; Ren, P.; Sagui, C. Classical Electrostatics for Biomolecular Simulations. *Chemical Reviews* **2014**, *114*, 779–814.

(38) Riniker, S. Fixed-Charge Atomistic Force Fields for Molecular Dynamics Simulations in the Condensed Phase: An Overview. *Journal of Chemical Information and Modeling* **2018**, *58*, 565–578.

(39) Alldritt, B.; Hapala, P.; Oinonen, N.; Urtev, F.; Krejci, O.; Canova, F. F.; Kannala, J.; Schulz, F.; Liljeroth, P.; Foster, A. S. Automated Structure Discovery in Atomic Force Microscopy. *Science Advances* **2020**, *6*, eaay6913.

(40) Grill, L.; Dyer, M.; Lafferentz, L.; Persson, M.; Peters, M. V.; Hecht, S. Nano-Architectures by Covalent Assembly of Molecular Building Blocks. *Nature Nanotechnology* **2007**, *2*, 687–691.

(41) Blunt, M. O.; Russell, J. C.; Gimenez-Lopez, M. d. C.; Garrahan, J. P.; Lin, X.; Schroder, M.; Champness, N. R.; Beton, P. H. Random Tiling and Topological Defects in a Two-Dimensional Molecular Network. *Science* **2008**, *322*, 1077–1081.

(42) Banhart, F.; Kotakoski, J.; Krasheninnikov, A. V. Structural Defects in Graphene. *ACS Nano* **2011**, *5*, 26–41.

(43) Setvin, M.; Wagner, M.; Schmid, M.; Parkinson, G. S.; Diebold, U. Surface Point Defects on Bulk Oxides: Atomically-Resolved Scanning Probe Microscopy. *Chemical Society Reviews* **2017**, *41*, 369.

(44) Coropceanu, V.; Cornil, J.; da Silva Filho, D. A.; Olivier, Y.; Silbey, R.; Brédas, J.-L. Charge Transport in Organic Semiconductors. *Chemical Reviews* **2007**, *107*, 926–952.

(45) Gordon, O. M.; Moriarty, P. J. Machine Learning at the (Sub)Atomic Scale: Next Generation Scanning Probe Microscopy. *Machine Learning: Science and Technology* **2020**, *1*, 023001.

(46) Butler, K. T.; Davies, D. W.; Cartwright, H.; Isayev, O.; Walsh, A. Machine Learning for Molecular and Materials Science. *Nature* **2018**, *559*, 547–555.

(47) Carleo, G.; Cirac, I.; Cranmer, K.; Daudet, L.; Schuld, M.; Tishby, N.; Vogt-Maranto, L.; Zdeborová, L. Machine Learning and the Physical Sciences. *Reviews of Modern Physics* **2019**, *91*, 045002.

(48) Himanen, L.; Geurts, A.; Foster, A. S.; Rinke, P. Data-Driven Materials Science: Status, Challenges, and Perspectives. *Advanced Science* **2019**, *6*, 1900808.

(49) Jesse, S.; Kalinin, S. V. Principal Component and Spatial Correlation Analysis of Spectroscopic-Imaging Data in Scanning Probe Microscopy. *Nanotechnology* **2009**, *20*, 085714.

(50) Woolley, R. A. J.; Stirling, J.; Radocea, A.; Krasnogor, N.; Moriarty, P. Automated Probe Microscopy *via* Evolutionary Optimization at the Atomic Scale. *Applied Physics Letters* **2011**, *98*, 253104.

(51) Kalinin, S. V.; Sumpter, B. G.; Archibald, R. K. Big–Deep–Smart Data in Imaging for Guiding Materials Design. *Nature Materials* **2015**, *14*, 973–980.

(52) Kalinin, S. V.; Strelcov, E.; Belianinov, A.; Somnath, S.; Vasudevan, R. K.; Lingerfelt, E. J.; Archibald, R. K.; Chen, C.; Proksch, R.; Laanait, N.; Jesse, S. Big, Deep, and Smart Data in Scanning Probe Microscopy. *ACS Nano* **2016**, *10*, 9068–9086.

(53) Rashidi, M.; Wolkow, R. A. Autonomous Scanning Probe Microscopy *in Situ* Tip Conditioning through Machine Learning. *ACS Nano* **2018**, *12*, 5185–5189.

(54) Gordon, O. M.; Hodgkinson, J. E. A.; Farley, S. M.; Hunsicker, E. L.; Moriarty, P. J. Automated Searching and Identification of Self-Organized Nanostructures. *Nano Letters* **2020**, *20*, 7688–7693.

(55) Leinen, P.; Esders, M.; Schütt, K. T.; Wagner, C.; Müller, K.-R.; Tautz, F. S. Autonomous Robotic Nanofabrication with Reinforcement Learning. *Science Advances* **2020**, *6*, eabb6987.

(56) Azuri, I.; Rosenhek-Goldian, I.; Regev-Rudzki, N.; Fantner, G.; Cohen, S. R. The Role of Convolutional Neural Networks in Scanning Probe Microscopy: A Review. *Beilstein Journal of Nanotechnology* **2021**, *12*, 878–901.

(57) Ronneberger, O.; Fischer, P.; Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation 2015, 1505.04597. arXiv, `http://arxiv.org/abs/1505.04597` (accessed 11/06/2021).

(58) Oktay, O.; Schlemper, J.; Folgoc, L. L.; Lee, M.; Heinrich, M.; Misawa, K.; Mori, K.; McDonagh, S.; Hammerla, N. Y.; Kainz, B.; Glocker, B.; Rueckert, D. Attention U-Net: Learning Where to Look for the Pancreas 2020, 1804.03999. arXiv, `http://arxiv.org/abs/1804.03999` (accessed 11/06/2021).

(59) Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Kopf, A.; Yang, E.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J., et al. In *Advances in Neural Information Processing Systems 32*, Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R., Eds.; Curran Associates, Inc.: New York, 2019, pp 8024–8035.

(60) Oinonen, N.; Xu, C.; Alldritt, B.; Canova, F. F.; Urtev, F.; Krejčí, O.; Kannala, J.; Liljeroth, P.; Foster, A. S. Electrostatic Discovery Atomic Force Microscopy. 2021, 2108.04333. arXiv, `http://arxiv.org/abs/2108.04333` (accessed 11/05/2021).

Supporting Information: Electrostatic Discovery Atomic Force Microscopy

Niko Oinonen,[1*], Chen Xu,[1*], Benjamin Alldritt,[1*]

Filippo Federici Canova,[1,2] Fedor Urtev,[1,3] Shuning Cai,[1] Ondřej Krejčí,[1]

Juho Kannala,[3] Peter Liljeroth[1†] and Adam S. Foster[1,4†]

[1]Department of Applied Physics, Aalto University, 00076 Aalto, Helsinki, Finland

[2]Nanolayers Research Computing Ltd, London N12 0HL, United Kingdom

[3]Department of Computer Science, Aalto University, 00076 Aalto, Helsinki, Finland

[4]WPI Nano Life Science Institute (WPI-NanoLSI), Kanazawa University,

Kakuma-machi, Kanazawa 920-1192, Japan

[*]These authors contributed equally.

[†]To whom correspondence should be addressed;

E-mail: peter.liljeroth@aalto.fi; adam.foster@aalto.fi.

Methods

Machine learning

The core of our model has the structure of the U-Net [1], where the feature maps first enter an encoder which down-samples them with pooling layers and then enter a decoder which up-samples them back to the original size, with skip connections between the layers of matching size in the encoder and decoder. In addition to the different number of channels and layers, the main difference to the original U-Net is that we start the network with 3D convolutions and then change to 2D convolutions in the middle, and we use Attention-Gate (AG) [2] layers in the skip connections. Around the core, we have the input stage which merges the two input sets of AFM images, and the output stage which outputs the ES Map descriptor. All of the convolutional layers use replicate padding, and, except for the output layer, all convolutional layers are followed by LeakyReLU activations [3] with negative slope of 0.1.

The model architecture, along with the shapes of the layers assuming 128×128 lateral input size, are illustrated in Fig. S1. At the start, the two inputs are fed into their own blocks of two 3D convolutions with 32 channels. The outputs from these are concatenated together in the channel dimension and then fed into the encoder. The encoder consists of three blocks of 3D convolutions and poolings. In each block, there are three 3D convolutions, with 48, 96, and 192 channels in each block, respectively, and the poolings are AvgPool layers. The poolings all have pool regions of $2 \times 2 \times 2$, but the middle pooling has a stride of $2 \times 2 \times 1$, so that the size of the feature map in z-direction is only reduced by 1. After the last pooling layer, the 3D feature maps are transformed into 2D feature maps by concatenating the remaining z-layers of the 3D feature maps into channels of the resulting 2D feature maps. The middle section between the encoder and the decoder has a block of three 2D convolutions with 512 channels. The decoder has three up-sampling stages corresponding to the three down-sampling stages of
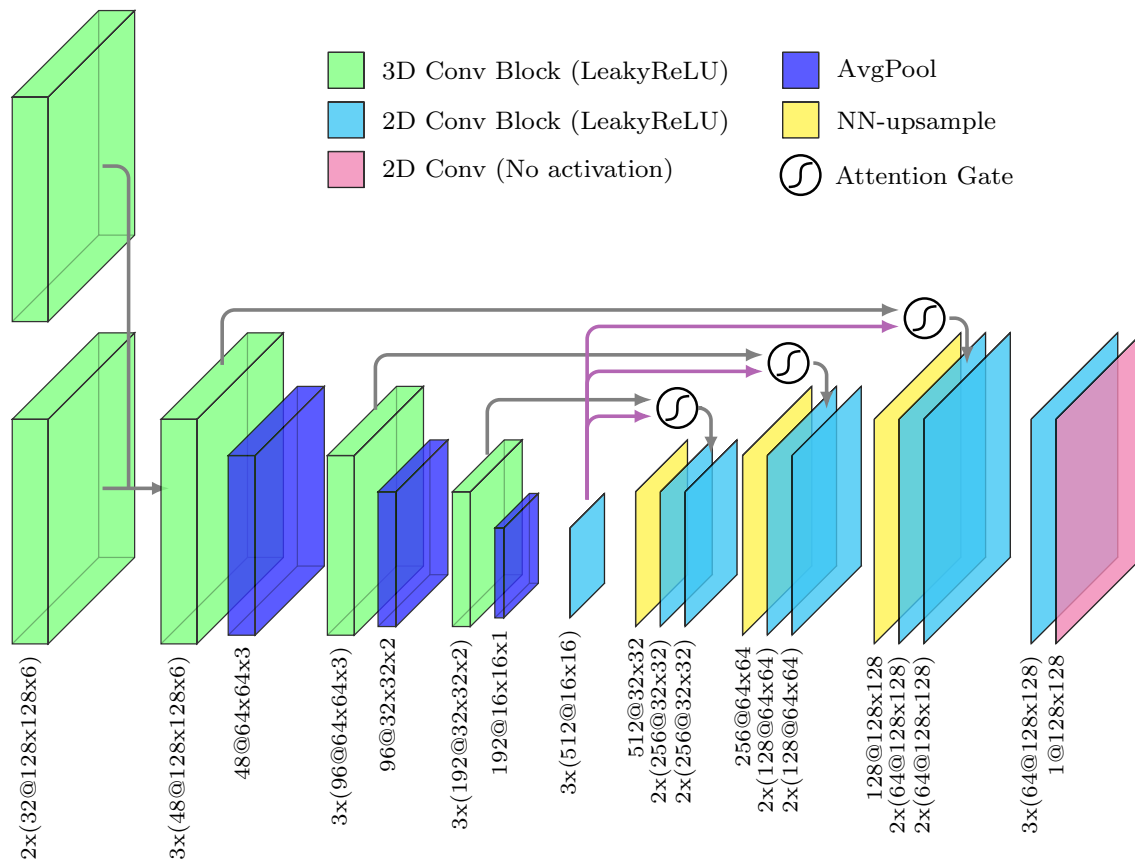
Figure S1: Schematic illustration of the model architecture. Below each layer or block of layers, the output shape of the layer is reported in the format (number of channels)@(feature map size) assuming a 128×128 lateral input size, and in the blocks the first number indicates the number of layers of that shape in the block. Here, NN = nearest neighbour.

the encoder. At each stage there is first a nearest-neighbour up-sampling followed by a block of two 2D convolutions. Then the skip connection from the corresponding stage of the encoder is passed through the AG and is concatenated as additional channels to the input of a second block of 2D convolutions. The convolution blocks in the decoder stages have 256, 128, and 64 channels. After the decoder, the model has one more 2D convolution block with three 2D convolutions and 64 channels, and one more 2D convolution with a single channel to output the ES Map descriptor. The total number of parameters in the model is 15,604,900.

In the proposed architecture we implemented AGs [2] on the skip connections from the

encoder to the decoder. Based on a specific task, AGs can suppress irrelevant and highlight useful parts in inputs. An AG architecture is illustrated in Fig. S2. It has two inputs: a set of feature maps from 3D convolution blocks in encoder flattened into 2D feature maps (x) and a query (q) – feature maps from the last 2D convolution layer in the middle part of the model. For a 128×128 lateral input size, the skip connections have the following shapes: 288@128x128, 288@64x64, 384@32x32 and the query shape is 512@16x16. After an interpolation of q to match the shape of x, both x and q are passed through independent 2D convolution layers with ReLU activations and then combined together by channel-wise summation. The feature maps are then passed through a 2D convolution with a Softmax activation to create a single-channel map of attention coefficients – the attention map. Finally, the attention map is mixed with the skip connection by element-wise multiplication in each channel. Due to the construction with a Softmax activation, the AG learns to highlight the most relevant regions in the input without explicitly being trained to do so.
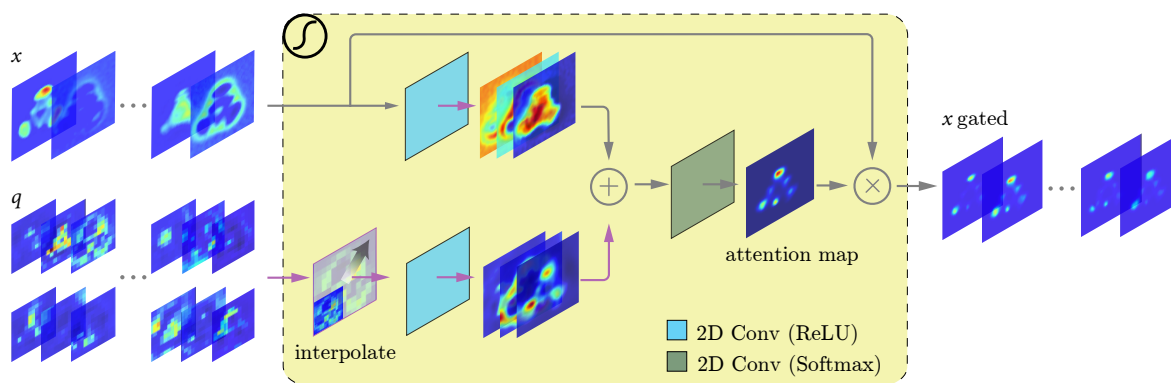

Figure S2: Schematic illustration of the Attention Gate (AG) using the BCB molecule as an example. Randomly picked features maps in inputs and outputs are presented. AG operates with 2 inputs: skip connection feature maps (input x) together with compressed representation at the end of the encoder (query q). Since all three skip connections have different sizes than the query, q is interpolated to match the size of x. Both x and q are passed through 2D convolutions with ReLU activation, and then they are summed together and the result is passed through a 2D convolution layer with Softmax activation to produce the attention map. The attention map is finally multiplied pixel-wise with the skip connection features maps to produce the gated output of the AG layer.

Our objective function is the mean squared error

$$\text{MSE}(y, \tilde{y}) = \frac{1}{N} \sum_{i=0}^{N-1} (y_i - \tilde{y}_i)^2, \tag{1}$$

where y is the predicted ES Map, $\tilde{y}$ is the reference ES Map, and the sum is over the N pixels. For reference, the losses on the final trained model are 2.17×10^{-5} on the training set, 2.49×10^{-5} on the validation set, and 2.47×10^{-5} on the test set. The parameters are optimized using the Adaptive Moment Estimation (Adam) optimizer [4], with learning rate 10^{-4} and the default values of $\beta_1 = 0.9$ and $\beta_2 = 0.999$ for the moment decay parameters. Additionally, we use a learning rate decay, where on each iteration i, the initial learning rate is multiplied by a factor

$$\frac{1}{1 + 10^{-5} \cdot i}. \tag{2}$$

The training set has a total of 6000 batches with 30 samples each, and the model is trained for a total of 50 epochs. The dataset is described in more detail below in Sec. "Dataset".

During training, we preprocess the samples in several ways before they enter the model. The samples are normalized by subtracting the mean and dividing by the standard deviation per each height-layer in the AFM image stack. For regularization, we randomly add to each sample noise, pixel shifts, cutouts, and additive background gradient planes, and the samples are randomly rotated, flipped, and cropped. The noise augmentation is discussed below in Sec. "Noise amplitude distributions". The pixel shifts are applied independently to each layer in the AFM image stack, such that the pixel values roll over the borders. The maximum shift between adjacent slices in the AFM image stack is 2% of the image size and maximum total shift is 4% of the image size. The cutouts randomly erase an area of the input image. The erased area for each cutout is at most 1% of the total area of the image and has a maximum aspect ratio of 1 to 10. A maximum of five cutouts are added to each image with 20% probability for each one. For details of the background gradient augmentation, see below Sec. "Surface tilt effect on model predictions". The original samples generated in a size of 192×192 are rotated to a random

angle by bicubic interpolation, flipped up-down with 50% probability, and then cropped to size 128×128 to get rid of any empty pixels in the corners. The images are then further cropped to a random position at random size of at minimum 75% of the original size and a random aspect ratio of at most 1.25 in either direction.

The training samples are all generated on a $24 \times 24 \, \mathring{A}^2$ frame discretized on a 192×192 grid, and the molecule is always in the center of the frame. Since the model is trained on this specific pixel density of $24 \, \mathring{A}/192 \, \mathrm{pixels} = 0.125 \, \mathring{A}/\mathrm{pixel}$, the experimental images are resized to match this resolution before entering the model. Additionally, we always crop the images into multiples of 8 pixels in each dimension in order to keep the dimensions consistent over the pooling and upsampling layers in the model (three halvings = 1/8 image size). The experimental images are also normalized in the same way as the simulated training samples.

Distance randomization

In AFM experiments, it is often difficult to know the exact distance between the tip and the sample, and the distance range where the tip-sample interaction is stable differs between samples. These facts mean that the range of distances available in AFM images is variable. In order to be robust against varying tip-sample distance, we randomize it during the generation of the training simulation samples within a $0.5 \, \mathring{A}$ window. Here we have to take into account the additional factor of the second tip. It is unclear whether the ML model would benefit from having the two tips at the same distance from the sample or if it would also work if the tips are not aligned.

In order to test this, we train the ML model with two differently generated datasets, one with matched tip distance where the tip-sample distance is the same for the two tips for the same training sample, and one with independently randomized tips. We then test how the MSE loss behaves for the two differently trained models as a function of the tip-sample distance on

a subset of 3000 samples from the test set (Fig. S3A,B). For matched tip distances on the test samples (Fig. S3A), we find that both models have almost a flat loss curve within the window of distances used in the training, but outside of that window the loss starts to increase. The increase in loss is especially sharp on the side of smaller distances. We discuss why too close distances are undesirable at more length in the context of the experimental predictions below in Sec. "Distance dependence". In these results the difference between the matched and independent tips is small, with possibly a small advantage in favor of the matched tips. However, when we do the test such that the CO-tip is held at constant height and the Xe tip is shifted, the difference in performance becomes very apparent. The model trained on matched tips does well for the zero-shift where the tip distances are matched, but when the Xe distance is varied the loss becomes significantly bigger, by more than a order magnitude even within the training window of distances. This is in contrast with the model trained with independently randomized tips, which has similarly flat loss curve as in the first test. Clearly, any small disadvantage for the independent tip randomization in the first test is worth the trade-off for significantly improved stability when the tip distances are not exactly matched.

Noise amplitude distributions

Experimental AFM images always have some level of noise present in the values of the pixels. In order to be robust against noise in the input images, we add noise with random uniform distribution to the simulated images during training. Since the noise is independent between training epochs, this also serves as a type of regularizing augmentation that reduces overfitting of the model. The generated noise is multiplied by the range of the values $(\mathrm{max} - \mathrm{min})$ in the sample before the addition operation to keep the level of the noise consistent between the samples. We test here three different ways of choosing the amplitude of the noise: constant amplitude, uniform random amplitude, and normally distributed amplitude. For the normal

Figure S3: Loss statistics with different randomizations of the tip-sample distance and the noise amplitude on a subset of the test set. (A, B) The MSE loss as a function of tip-sample distance offset dh with (A) both tips offset and (B) only Xe offset for two models trained with independently randomized and jointly randomized distance for the two tips. Here $dh = 0\,\text{Å}$ represents the average distance used in the training. (C) The MSE loss as a function of noise amplitude for three different models trained with constant ($\mathcal{C}$), uniform random ($\mathcal{U}$), and normally distributed ($\mathcal{N}$) noise amplitudes. In all plots, the solid lines represent the mean loss, and the dashed lines represent the 5th and 95th percentiles, so that 90% of the losses are contained within the region enclosed by the dashed lines of the same color.

distribution, we use the absolute value of the generated value as the amplitude, and we choose the standard deviation of the normal distribution to be 0.1. This gives the noise amplitude an expected value of ~ 0.08. To keep the average level of the noise consistent between the tests, for the constant amplitude we choose the value 0.08, and for the uniform random amplitude we choose the range $[0, 0.16]$.

We test these three differently trained models on a subset of 3000 samples from the test set. The average MSE loss on these test samples as a function of the noise amplitude is presented in Fig. S3C. The most striking feature here is the difference between the constant amplitude model and the random amplitude models. The model trained with the constant amplitude does the best on the amplitude of the noise that it was trained on and has worse loss for all other amplitudes, including the zero-amplitude without any noise. This is saying that for this model clean images are harder to interpret than noisy ones, clearly an undesired behaviour. In contrast, the models

trained with random noise amplitude have much flatter loss curves, with the uniform random amplitude having a small advantage over the normally distributed one. Further tests would be needed to determine what is the optimal distribution for the random noise amplitude, but it is clear that random amplitude for the noise is better than constant amplitude. For the training of the model used for the predictions in the main article, we used the normally distributed amplitude.

Dataset

Our model training is based on a database of 81086 molecular geometries containing the elements H, C, N, O, F, Si, P, S, Cl, and Br. The distribution of the elements in the molecules is shown in Table T1. Here we can see that the distribution is not even: H and C are contained in almost every molecule with N and O being very common as well, but the rest of the elements are significantly less common. In our previous work on molecule structural prediction from AFM images we used a simple criterion with a fixed number of rotations for each molecule to choose the molecule orientations for the samples [5]. This lead to an overemphasis on the more common elements, especially H, in the dataset. Here we have chosen the rotations for the molecules more carefully in order to make the element distribution more even in the dataset.

For choosing the rotations, we want to consider what elements are close to the surface of the molecule, so that those atoms could possibly be seen in the AFM images. To this end, we compute the convex hull of the molecule, yielding us sets of three points that define planes on the surface of the molecule. We consider each one of these planes in turn and include the rotation corresponding to the plane probabilistically based on the elements close to the plane, choosing the probabilities such that the rarer elements are emphasized. An element is considered to be close to the plane if an atom with that element is within $0.7\,\text{Å}$ of the plane. To counter any bias that using the convex hull planes may incur, we also choose completely random rotations of the

molecules, which are again included probabilistically emphasizing the rarer elements. Finally, we noted that the database does not contain many completely flat geometries, so we include any rotations of the molecules that contain a planar segment, which we define as a plane on the surface of the molecule which contains at least 10 atoms within $0.1\,\text{Å}$ of the plane. In order not to have overlap between the rotations, no rotations within $5°$ of each other for the same molecule are included.

Using this procedure, we generate a total of 235554 different orientations of the molecules, which we divide into training, validation, and test sets as 180000, 20000, and 35554 samples, respectively. We take care not to include any of our test molecules in the training or validation sets. The distribution of the elements contained in the final chosen rotations based on the $0.7\,\text{Å}$ criterion is shown in Table T1. H and C are still the most common elements, and this is natural, since any orientation where one of other elements is seen, likely H and/or C is also there. The occurrence of the rest of the elements is now more even, except for Si and P, which are mostly only contained inside the molecules and therefore are not often seen close to the surface.

Element	% of molecules in database	% of chosen rotations
H	99.3	87.3
C	99.6	49.8
N	60.8	24.8
O	76.5	29.5
F	5.4	16.5
Si	1.5	0.2
P	3.1	1.2
S	15.0	23.4
Cl	13.7	28.6
Br	3.1	16.4

Table T1: Distribution of different elements contained in the molecules in our database and in the rotations of the molecules that we chose. For the chosen rotations an element is included in the count if it is contained in the region up to $0.7\,\text{Å}$ below the top-most atom in the molecule.

ES Map descriptor

The ES Map descriptor is the z-component of the ES field originating from the charges in the sample molecule, calculated at a constant-height surface $4\,\text{Å}$ above the top-most atom in the molecule, and then cut to be non-zero only in the region occupied by the molecule. This process is illustrated in Fig. S4. We define the z-direction to be parallel to the oscillation direction of the AFM probe, which is perpendicular to the hypothetical surface which the molecule would be sitting on, and the positive z-direction is pointing away from the surface (out of the page in the figures here). We find the highest z-coordinate of a center of an atom in the molecule and then add $4\,\text{Å}$ to that value to define the z-coordinate of the constant-height surface of the ES Map. The xy-coordinates of the pixels form a grid corresponding to the matching pixel coordinates in the AFM images. Then for a given pixel $\mathbf{R}_{ij} = (x_i, y_j, z)$ the value of the pixel is

$$E_{\mathrm{z}}(\mathbf{R}_{ij}) = k_{\mathrm{e}} \sum_{k=1}^{n} \frac{q_k (\mathbf{R}_{ij} - \mathbf{r}_k) \cdot \hat{\mathbf{z}}}{|\mathbf{R}_{ij} - \mathbf{r}_k|^3}, \tag{3}$$

where k_{e} is the Coulomb constant, q_k is the charge and $\mathbf{r}_k$ is the coordinate of the kth atom in the molecule, n is the number of atoms in the molecule, and $\hat{\mathbf{z}}$ is a unit vector in the z-direction. For restricting the non-zero area, we use the vdW-Spheres descriptor, which we introduced in our previous work [5]. We modify the descriptor here by adding a constant $1\,\text{Å}$ to the vdW radii of the atoms and restrict the deepest coordinate to be $2\,\text{Å}$ below the top coordinate. We then turn the vdW-Spheres descriptor into a binary mask by setting the background values to 0 and all other values to 1. This mask is then multiplied pixel-wise with the E_{z}-values computed earlier to produce the final pixel values of the ES Map descriptor.

AFM Simulations

We simulate AFM images using the probe particle model [6]. The procedure for generating the training samples is explained in our previous work [5]. Here we additionally do simulations
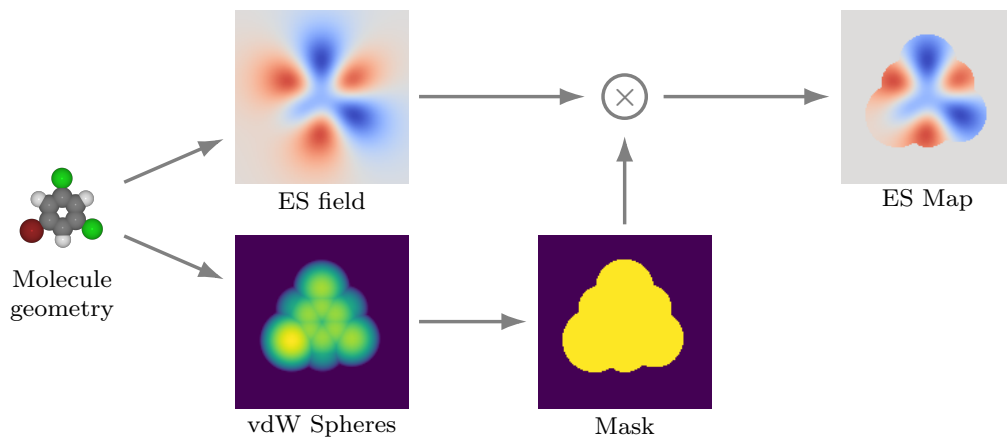
Figure S4: Schematic illustration of the process for generating the ES Map descriptor using the BCB molecule as an example. The molecule geometry is used to compute both the ES field and the vdW-spheres descriptor. The vdW-Spheres descriptor gets turned into a binary mask which is then multiplied pixel-wise with the ES field to produce the ES Map descriptor.

with the Xe and Cl probe-particle tips. The lateral spring constants we use for the probe particles are $0.25\,\text{N/m}$ for CO and Xe, and $0.5\,\text{N/m}$ for Cl. The radial spring constant is $30\,\text{N/m}$ in all cases. The CO and Xe tip charges are modelled as quadrupoles with quadrupole moments of $-0.1\,\text{e} \times \text{Å}^2$ and $0.3\,\text{e} \times \text{Å}^2$, respectively, and the Cl tip charge is modelled as a monopole with a charge of $-0.3\,\text{e}$, where e is the elementary charge. The Lennard-Jones parameters for each atom type contained in our database of molecules are listed in Table T2. To regularize the model and make it more robust, we randomize the tip-sample distance in the simulations within a $0.5\,\text{Å}$ window (see Sec. "Distance randomization" above for more details). Additionally, the lateral equilibrium position of the probe particle is randomized within a disk of radius $0.5\,\text{Å}$.

Sensitivity of predictions to spring constant values

We use fixed values for the lateral and radial spring constants in the AFM simulations in the training set. Since the tip condition can vary between AFM experiments, it is worth considering how sensitive the simulation and the predictions are to the chosen spring constant values. To this end, we run simulations on a subset of the test set varying the spring constant values in

a range of $0.20\ldots0.30$ N/m for the lateral spring constant k_{lat}, and $20\ldots40$ N/m for the radial spring constant k_{rad}. On visual inspection of the simulated AFM images, for the radial spring constant there is no discernible difference between the different values in the chosen range. For the lateral spring constant, the differences are small, but can still be observed as a gradual change to a slightly sharper contrast in the close range with higher values of the spring constant. To quantify the sensitivity of the model predictions, we run the predictions for the simulated images and record the MSE loss as a function of the spring constant values (Fig. S5). The result is in line with the visual inspection of the AFM images: for the radial spring constant there is no significant difference in the loss values with different k_{rad} values, and for the lateral spring constant the loss increases smoothly when deviating from k_{lat} value used in the training. The loss increases more with decreasing k_{lat}, reaching a value roughly 3 times the minimum loss at $k_{\mathrm{lat}} = 0.25\,\mathrm{N/m}$. Keeping in mind that the MSE loss emphasizes outliers, this does still not correspond to a very large decrease in average performance. However, the result does show that the lateral spring constant is a parameter that could be worth randomizing during the training in the future to be more robust against changes in the tip condition.

Element	$R_{ii}[\text{Å}]$	$E_{ii}[\text{eV}]$
H	1.4870	0.000681
C	1.9080	0.003729
N	1.7800	0.007372
O	1.6612	0.009106
F	1.7500	0.002645
Si	1.9000	0.025490
P	2.1000	0.008673
S	2.0000	0.010841
Cl	1.9480	0.011491
Br	2.2200	0.013876
Xe	2.1815	0.024344

Table T2: Lennard-Jones parameters used in the probe particle simulations.

Figure S5: Loss statistics for different values of the spring constants used in the AFM simulation for a subset of the test set. The plots show the MSE loss as a function of the value of (A) the lateral spring constant k_{lat} and (B) the radial spring constant k_{rad}. The spring constants are altered for both the CO and the Xe tips. In both plots, the solid lines represent the mean loss, and the dashed lines represent the 5th and 95th percentiles, so that 90% of the losses are contained within the region enclosed by the dashed lines.

Experimental

The AFM images were taken on a combined non-contact AFM/STM system (CreaTec) with a commercial qPlus sensor with a Pt/Ir tip, operating at T $\approx$ 5K in ultrahigh vacuum at a pressure of ~1 × 10⁻¹⁰ mbar. The qPlus sensor had a resonance frequency of $f_0 \approx$ 30046 Hz, a quality factor Q = 67714, and was always operating with an oscillation amplitude of A = 50 pm.

The Cu(111) substrate (MaTeck) was prepared by repeated Ne⁺ sputtering with a beam energy of 750 eV and ion current of 20 μA for 15 min followed by annealing at 520~550°C for 5 min. A flat Cu (111) surface with large terrace and minimum amount of impurities was often obtained within 3 cycles. The 1-Bromo-3,5-dichlorobenzene molecules (Sigma-Aldrich; purity 98%) were deposited onto the substrate at ~5 K through a variable leak valve **1** at a chamber pressure of 1 × 10⁻⁶ mbar for 30 seconds. Then the CO molecules (Praxair; purity 99.997%)

Figure S6: Full sets of experimental AFM images for (A) BCB (CO), (B) BCB (Xe), (C) PTCDA (CO), and (D) PTCDA (Xe). The scale bars are 5 Å long.

15

Figure S6: (Continued) Full sets of experimental AFM images for (E) Water (CO), (F) Water (Xe). The scale bars are 5 Å long.

were deposited onto the substrate at 5 K through a variable leak valve **2** at a chamber pressure of 1×10^{-6} mbar for less than 5 seconds. Finally, the Xe atoms (Fluka; purity 99.995%) were deposited onto the substrate at 5 K through the variable leak valve **1** at a chamber pressure of 1×10^{-6} mbar for 30 seconds.

Tip conditioning was usually performed by controlled contact with the Cu substrate and/or by applying a 1 second voltage pulse of 3~10 V, both with feedback turned off. The tip was deemed as good when a symmetric contact mark was observed as well as a reasonably resolution of the organic molecules was achieved. The tip apex was believed to be covered with Cu atoms

after these operations.

The constant height AFM images were taken with metal tips functionalized with a single CO molecule or a single Xe atom. The CO functionalization was achieved by applying a set-point of 8 mV /100 pA with the tip over a CO molecule, followed by turning off the feedback and then ramping the sample bias from zero to 2.6 V. A sudden decrease in the current happened at about 2.2 V indicates a successful functionalization. A subsequent scan over another CO showing sharp central protrusion can confirm the functionalization. After finishing with the CO tip, a bias ramping from zero to 3.6 V with feedback turned off can remove the CO while minimizing the perturbation to the structure of the metal tip apex. A sudden change of current at around 3.2 V often indicates a successful removal of the CO.

A second sequence of AFM images of the same molecule was taken with a Xe functional-ized tip. The Xe functionalization was achieved by applying a set-point of 100 mV / 100 pA, followed by turning off the feedback and then bringing the tip into contact with a cluster of Xe atoms. A sudden decrease in current happened at ~3.5 Å advanced from the starting position indicates a successful transfer of a Xe atom. An STM scan with set-point of 100 mV / 100 pA capable of resolving individual Xe atoms inside the Xe cluster can further confirm such a functionalization.

The second experiment with PTCDA molecules (Sigma-Aldrich; purity 97%) was done in a similar manner, except the PTCDA molecules were deposited onto the Cu(111) substrate at about 200 K using thermal sublimation.

The third experiment was done with water molecules (Sigma-Aldrich SKU38796; deion-ized). The water was purified before deposition, it was firstly boiled at 100°C to rid of any residual gas inside, and was then degassed thoroughly *via* several freeze-pump-thaw cycles. During the experiment, water molecules were deposited *via* a variable leak valve **3** aiming di-rectly at the Cu(111) inside the scanner held at 5 K. The sample was subsequently heated up to

17

40 K, so that water molecules started to form clusters [7]. The sample was cooled back to 5 K thereafter. Xe and CO were deposited onto the surface with the same procedure as before.

Extended results

Figure S7: Predictions for the benchmark examples using the DFT Hartree potential for electrostatics in the simulations. Compare to Fig. 2 in the main article.

Single-channel measurements

Since the two-tip measurement presents an additional experimental challenge, we also try training a model using only a single-tip input of CO-AFM. This model is the same as the two-tip model, except that it lacks the other branch of layers in the beginning of the network. Fig. S9A

Figure S8: Prediction and reference for on-surface geometry of PTCDA using the DFT Hartree potential for electrostatics in the AFM simulations and for the reference ES Map descriptor.


Figure S9: Single-tip characterization of the benchmark examples in the main paper. Predictions are shown for (A) simulated data of TTF-TDZ and experimental data of (B) BCB and (C) PTCDA.

shows an example prediction on simulated data of the TTF-TDZ molecule. At a glance the prediction matches really well with the reference, but a closer inspection reveals that the magnitude of the field is not as accurate. The relative error for the prediction is 5.56%, more than twice the value for the two tip model. When measured on the whole test set, the average loss for the single-tip model is 77% higher than for the two-tip model. Therefore, the single-tip model is less robust, but performance is not fatally worse.

We also apply the single-tip model to experimental data of BCB and PTCDA (Fig. S9B,C) and find that the predictions are not very sensible. For BCB the model predicts mostly negative charge over the whole molecule with some positive regions over two of the halides. This does not match with either of the reference descriptors, where we expect to find the halides to be the most negative regions. The prediction for the PTCDA is similarly biased towards negative values in the middle of the molecule which in both reference descriptors is the least negative

region. These results show that currently the addition of second channel of information is necessary to make the prediction work. Still, the relatively good performance on the simulations indicates that if the simulation model could be improved to be more accurate, then possibly even a single-channel measurement could be used for prediction.

Other tip combinations

Figure S10: Predictions with models trained on Cl-CO and Cl-Xe tip combinations on the benchmark examples. Compare to Fig. 2 in the main article.

In order to show that the specific combination of CO and Xe tips is not special, we also generate simulations with the Cl tip and train models using the alternative tip combinations of Cl-CO and Cl-Xe. Figure S10 shows example predictions for both tip combinations on simulations of the three benchmark examples introduced in the main article. On these examples, we find that the performance is roughly on par with the CO-Xe model, and for the losses on the test set we even find that the Cl-CO and Cl-Xe models have lower average losses than the CO-Xe model, by 46% and 43%, respectively. In principle, any combination of tips can be used, as long as accurate simulated training data can be generated for them.

Distance dependence

The model is trained to take in AFM image stacks with six constant-height slices for both tips. In training the model, we consciously choose the tip-sample distance to be in range where the furthest images are mostly in the attractive regime, where only the overall shape of the molecule can be distinguished, and the closest images are in the repulsive regime, where at least some sharp atomic features are seen. However, we do not want to go too close to the molecule for two reasons. First, at close range the simulation data used for the training is less representative of the experimental case due to the simulation not taking into account any tip-induced relaxation of the sample. Second, at very close range the interaction between the tip and the sample is dominated by Pauli repulsion and the role of the electrostatics decreases.

In the experiments we have more than six slices for each measurement (see Figs. S6 and S6), which leaves us with some room to choose which subset of images we use for the prediction. This selection process is still not automated and we have to use some judgment in choosing what is the best range for the data so that it best matches the training data, though we do augment the training with a variable range of distance in an attempt to be robust against variations in the distance. In Fig. S11 we explore for two of our experimental cases, BCB and PTCDA, what
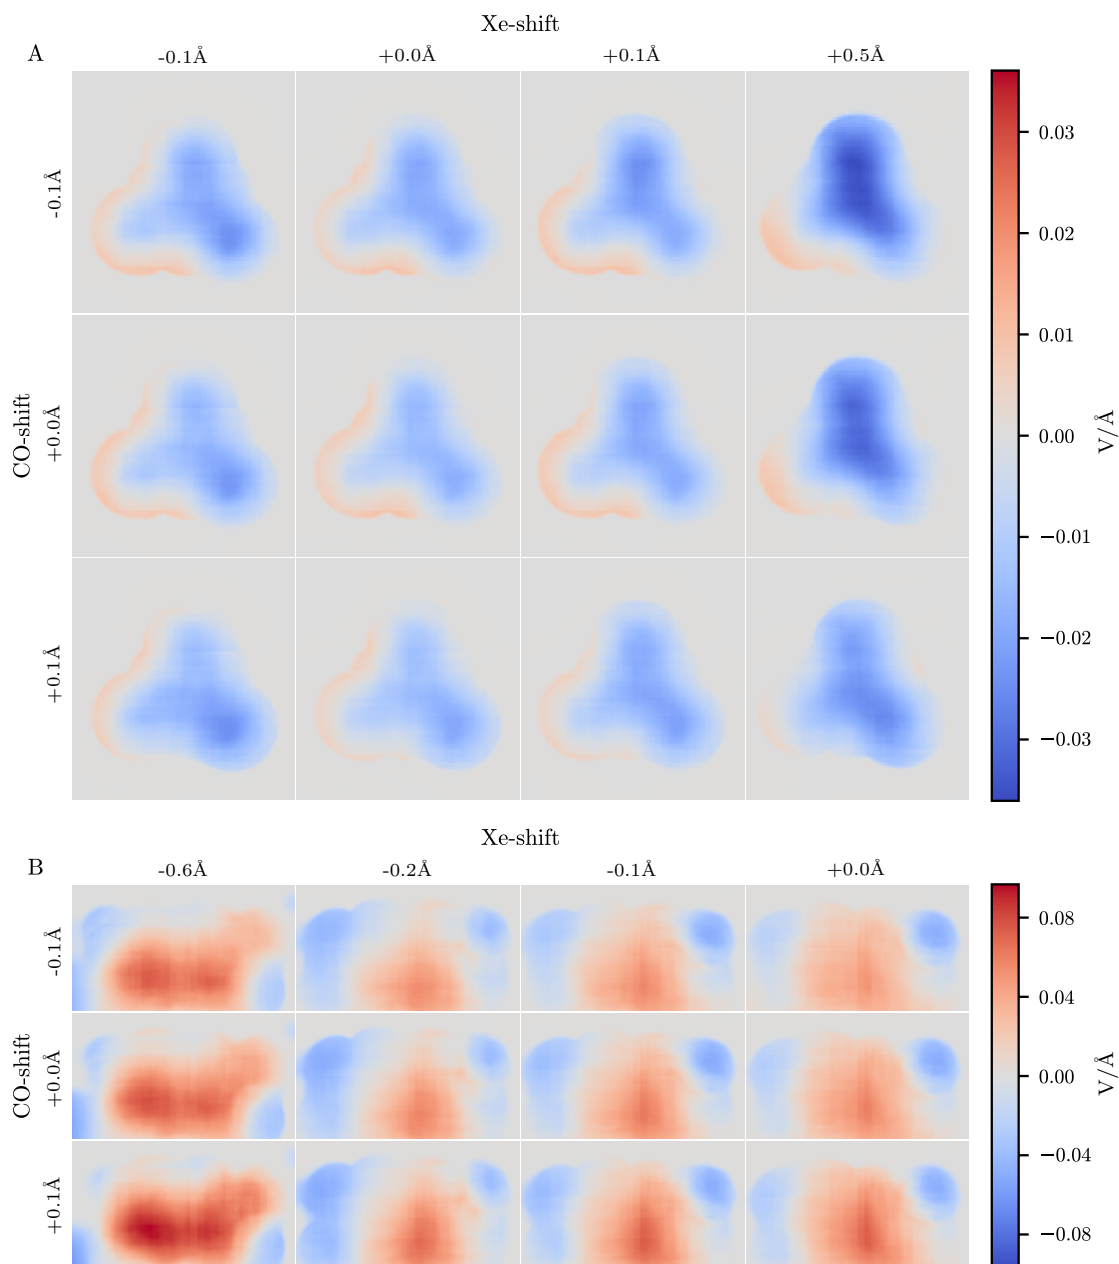
Figure S11: Predictions at different tip-sample distances for experimental images of (A) BCB and (B) PTCDA. On each row the CO-AFM input has been shifted closer or further and on columns correspondingly the Xe-AFM input has been shifted. Here, the shift of +0.0 Å corresponds to the distance used in the predictions in Figs. 3 and 4 of the main article. Negative shift corresponds to smaller tip-sample distance.

happens to the predictions when we vary the tip-sample distance of either the CO or the Xe tip. In both cases we find that small deviations within a range of $0.2\,\text{Å}$ do not change the predictions significantly. We also try larger deviations of $+0.5\,\text{Å}$ for BCB and $-0.6\,\text{Å}$ for PTCDA of the Xe tip to show that too large deviations start to alter the predictions.

Possible extra electron in PTCDA

Figure S12: AFM simulations of the on-surface PTCDA with one extra electron added to the top right oxygen. Note that the geometry has been flipped left-to-right compared to Fig. S8A.

In our experimental Xe-AFM images of PTCDA (Fig. S6D) we noted that there appears an unusual bright feature over the oxygen at the top right of the images. We speculate that this feature could be due to an extra electron trapped on that oxygen based on simulations with such an extra electron resulting in a similar bright feature, shown in Fig. S12. In this simulation we took the on-surface geometry used for the DFT Hartree simulation in Fig. S8A and did the simulation using Mulliken charges but adding an extra charge of $-1\,\text{e}$ to the oxygen at the top right of the molecule. The result is a large bright spot surrounded by a dark halo in the Xe-AFM image, somewhat similar to the one in the experimental images. Furthermore, in the CO-AFM simulation there appears a sharp change in contrast over the same oxygen, which on the closer distance makes the oxygen appear very dim compared to the rest of the molecule, which also bears similarity to the experimental images (Fig. S6C). It is, however, unclear if such extra

electron would stay trapped on the oxygen and not transfer to the substrate or tip even when the molecule is being pushed by the AFM probe.

## Surface tilt effect on model predictions

Figure S13: Surface tilt effect on model predictions. (A) ES Map prediction of experimental AFM images of PTCDA on a model trained without background gradient augmentation. (B) Simulated AFM images of PTCDA with added background gradient. Using the input data in (B) we predict the ES Map on models trained (C) without and (D) with the background gradient augmentation.

The PTCDA dataset presented us with the additional challenge that the experiment was performed at a slight tilt which resulted in a gradient in the background of the image. When we artificially added a similar gradient to the simulated AFM images of PTCDA (Fig. S13B), the prediction (Fig. S13C) failed by incorrectly predicting the positive region in the middle of the molecule as being close to neutral. The prediction of the experimental data (Fig. S13A) showed a similar pattern.

Motivated by this finding, we augmented the training of the model with these background gradients, implemented by adding a plane with a set gradient to each AFM image set. The

direction of the gradient is uniform random, and the magnitude of the gradient is randomized such that the range of values in the plane is at most 30% of the range of values in the image set. The zero-point of the plane is always at the center of the images. With this augmentation, we found that the predictions improved both on the simulated images (Fig. S13D) and the experimental prediction also became more consistent. Being robust against these kinds of tilts could be useful in situations where a tilted planar section of a molecule needs to be characterized [8].

References

(1) Ronneberger, O.; Fischer, P.; Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation 2015, 1505.04597. arXiv, `http://arxiv.org/abs/1505.04597` (accessed 11/06/2021).

(2) Oktay, O.; Schlemper, J.; Folgoc, L. L.; Lee, M.; Heinrich, M.; Misawa, K.; Mori, K.; McDonagh, S.; Hammerla, N. Y.; Kainz, B.; Glocker, B.; Rueckert, D. Attention U-Net: Learning Where to Look for the Pancreas 2020, 1804.03999. arXiv, `http://arxiv.org/abs/1804.03999` (accessed 11/06/2021).

(3) Xu, B.; Wang, N.; Chen, T.; Li, M. Empirical Evaluation of Rectified Activations in Convolutional Network 2015, 1505.00853. arXiv, `http://arxiv.org/abs/1505.00853` (accessed 11/06/2021).

(4) Kingma, D. P.; Ba, J. Adam: A Method for Stochastic Optimization 2017, 1412.6980. arXiv, `http://arxiv.org/abs/1412.6980` (accessed 11/06/2021).

(5) Alldritt, B.; Hapala, P.; Oinonen, N.; Urtev, F.; Krejci, O.; Canova, F. F.; Kannala, J.; Schulz, F.; Liljeroth, P.; Foster, A. S. Automated Structure Discovery in Atomic Force Microscopy. *Science Advances* **2020**, *6*, eaay6913.

(6) Hapala, P.; Kichin, G.; Wagner, C.; Tautz, F. S.; Temirov, R.; Jelínek, P. Mechanism of High-Resolution STM/AFM Imaging with Functionalized Tips. *Physical Review B* **2014**, *90*, 085421.

(7) Liriano, M. L.; Gattinoni, C.; Lewis, E. A.; Murphy, C. J.; Sykes, E. C. H.; Michaelides, A. Water–Ice Analogues of Polycyclic Aromatic Hydrocarbons: Water Nanoclusters on Cu(111). *Journal of the American Chemical Society* **2017**, *139*, 6403–6410.

(8) Albrecht, F.; Pavliček, N.; Herranz-Lancho, C.; Ruben, M.; Repp, J. Characterization of a Surface Reaction by Means of Atomic Force Microscopy. *Journal of the American Chemical Society* **2015**, *137*, 7424–7428.