

ONGOING DEVELOPMENTS ON CONTINUUM SOLVATION MODELS

Matthew Truscott

Dissertation Prepared for the Degree of

DOCTOR OF PHILOSOPHY

UNIVERSITY OF NORTH TEXAS

May 2022

APPROVED:

Oliviero Andreussi, Major Professor

Gerardo Andrés Cisneros, Committee  
Member

Marco Buongiorno Nardelli, Committee  
Member

Yuanxi Wang, Committee Member

Jingbiao Cui, Chair of the Department of  
Physics

Pamela Padilla, Dean of the College of  
Science

Victor Prybutok, Dean of the Toulouse  
Graduate School

Truscott, Matthew. *Ongoing Developments on Continuum Solvation Models*.

Doctor of Philosophy (Physics), May 2022, 154 pp., 3 tables, 34 figures, 2 appendices, 106 numbered references.

This work explores a continuum representation for diffuse layer models, thereby endowing continuum embedding models the ability to capture electrostatic phenomena in the environment such as the existence of electrolyte ions, and the nature of ionic liquids. It introduces a new field-aware continuum model that adjusts the size of the quantum regime per atom based on the distribution of charge in a system. The model accounts for the asymmetric nature of solvent distribution when applied to cations versus anions; it also overcomes the need to parameterize continuum interface models for different charged systems. The continuum representation of cavitation in water does not account for the tendency for water to form a hydrogen bonding network that is broken due to the formation of cavities. This effect is a major contributor to hydrophobic solvation and is an important precondition to the investigation of solvated proteins with continuum embedding. A new model inspired by machine learning advances is trained on molecular dynamics simulations due to the difficulty of isolating the cavitation energy term in experiment. Thermodynamic integration is used to calculate the energy from a step-like repulsive potential from cavities in TIP4P water, cavities ranging from small organic molecules, to small proteins. Predictions from this new model show a small improvement for small molecules and scale much better with respect to the size of the system.

Copyright 2022  
by  
Matthew Truscott

## TABLE OF CONTENTS

	Page
LIST OF TABLES	vi
LIST OF FIGURES	vii
CHAPTER 1 INTRODUCTION	1
1.1. Materials Discovery	1
1.2. Simulation Models	4
1.3. Continuum Solvation Specifically	7
CHAPTER 2 CONTINUUM SOLVATION	10
2.1. Polarized Continuum Model	11
2.2. Self Consistent Continuum Solvation Model	12
2.3. Soft-sphere continuum solvation	16
2.4. Model parameterization	17
2.5. Non-local models	18
2.6. Additional Energy Terms	20
CHAPTER 3 SIMULATION SOFTWARE	21
3.1. Quantum ESPRESSO	21
3.2. Environ	21
3.3. Pythonization	23
3.4. DL POLY	25
3.5. AiiDA	25
3.5.1. Solvation	27
3.5.2. Grand Canonical	27
3.5.3. Machine Learning Adsorption Energy	30
3.6. Visualization	30
CHAPTER 4 CONTINUUM MODELS OF THE ELECTROCHEMICAL DIFFUSE	

LAYER IN ELECTRONIC-STRUCTURE CALCULATIONS	33
4.1. Introduction	33
4.2. Methods	36
4.2.1. The Electrolyte Cavity	36
4.2.2. The Electrolyte Models	38
4.2.3. Computational Implementations	47
4.2.4. Computational Details	54
4.3. Results and Discussion	57
4.3.1. Vacuum	57
4.3.2. Implicit Solvent	65
4.4. Summary and Conclusions	76
4.5. Kohn-Sham Potential and Force Contributions	77
4.6. Analytical solution of the Poisson-Boltzmann model of the diffuse layer	78
4.7. Analytical solution of the linearized Poisson-Boltzmann model of the diffuse layer	80
 CHAPTER 5 FIELD-AWARE INTERFACES IN CONTINUUM SOLVATION	 81
5.1. Introduction	81
5.2. Functional Derivatives	84
5.3. Field-aware function	87
5.4. Computational Details	89
5.5. Results	91
5.5.1. Propanoic Acid	93
5.5.2. Diethyl Ether	95
5.5.3. Benzyl Alcohol	97
5.5.4. Water	99
5.6. Conclusions	102
 CHAPTER 6 MACHINE LEARNING	 103

CHAPTER 7 MACHINE LEARNING CAVITATION ENERGY	109
7.1. Scaled Particle Theory	111
7.2. Methods	115
7.2.1. General Framework	115
7.2.2. Descriptors	116
7.2.3. Molecular Dynamics	117
7.2.4. Machine Learning	121
7.3. Computational Details	121
7.3.1. Small Molecules	122
7.3.2. Proteins	123
7.4. Results	123
7.4.1. Convergence Testing for Molecular Dynamics	123
7.4.2. Cavitation Energy of Spheres	126
7.4.3. Linear Regression on Descriptors	130
APPENDIX A FIELD AWARE DERIVATIVES	135
APPENDIX B ERROR ANALYSIS	139
REFERENCES	143

## LIST OF TABLES

		Page
5.1	Best Parameter Set	91
7.1	The performance (given in terms of the mean absolute error with respect to molecular dynamics simulations) of scaled particle theory as optimized for small molecules (SM) is shown. These optimized parameters $y$ , $\alpha$ are then applied to a set of proteins (P).	129
7.2	All values above are MAE (kcal/mol). A subset of descriptors was used for the Environ surface and Environ volume descriptors, obtained using CUR decomposition. Best performing results are obtained for each combination of descriptors for the small molecule set and then applied to proteins. Best performing results are obtained from a mix of small molecules and proteins towards proteins and then applied to small molecules.	133

## LIST OF FIGURES

	Page
3.1	Workflow of a typical calculation in Environ. In order to implement new interface definitions, functional derivatives for the energy and force terms need to be derived. 22
3.2	Solvation workflow implemented in AiiDA consisting of two PW calculations, one in vacuum and one in solution. 27
3.3	Grand canonical simulation workflow consisting of pre- and post-processing stages, and a series of monolayer and bulk simulations. 29
3.4	Graph based workflow that performs runs a minimal series of simulations to generate training data for adsorbate configurations. 31
4.1	Visualization of the atomistic details of a typical electrochemical setup (A): The metal slab (silver spheres) is in contact with an neutral aqueous solution (oxygen in red, hydrogen in white) containing electrolyte species (cyan and yellow spheres). Continuum models are obtained by integrating out the atomistic degrees of freedom of the mobile species (e.g. water and electrolytes) and replacing them with homogeneous continuum bodies, whose boundaries reflect the physical separation between the QM system and the environment. In B the solvent boundary (green transparent surface) is reported together with the continuous charge density of the electrolyte (blue transparent field), as computed for a charged substrate. Different onset for the solvent and electrolyte continuum models can be imposed in the definition of the model or can be obtained by including additional repulsive interactions between the continuum electrolyte and the QM substrate. 34
4.2	2D-maps of the electrostatic potential computed with the planar-averaged analytical PB model (A) and the corresponding numerical version (B) for a Ag(100) slab in vacuum with a total charge $q = 0.028$ a.u.. The



planar-averaged electrostatic potential computed for the neutral and the charged surface (red and blue, respectively) is plotted as a function of the surface normal direction  $x$  in C and D. The same data is plotted in the C and D panels, using different scales for the potential axis. The planar interface with a symmetric monovalent electrolyte with bulk concentration  $c_0 = 0.1$  M has been employed.

55

4.3 The charge per surface atom (in a.u.) is plotted as a function of the potential (in V). The Helmholtz model has been used for all data sets, varying the  $d$  and  $\Delta$  parameters that define the position and the width of the countercharge density, respectively. Note that the three lines corresponding to  $d = d_0/2 + 6$  Bohr are essentially superimposed ( $d_0$  is the slab thickness). The inset shows the differential capacitance as a function of the potential, as calculated from the analytical derivative of the spline-interpolated charge-potential curves (same line styles as in the main plot).

58

4.4 Same as Figure 4.3, but for the planar-averaged analytic linearized-PB model (triangles and dashed lines) and the corresponding numerical implementation (circles and solid lines). Red, blue and green symbols correspond to bulk electrolyte concentrations  $c_0 = 1$  M,  $c_0 = 0.1$  M and  $c_0 = 0.01$  M, respectively.  $x_{Stern}$  in the analytic model (Section 4.2.3.2) and  $d$  in the planar interface employed in the numerical model (Section 4.2.3.3) are set so that the interface lies in both cases 6.568 Bohr away from the surface.

59

4.5 Same as Figure 4.4, but for the full PB model, in the planar-averaged analytic implementation (triangles and dashed lines) and the numerical implementation (circles and solid lines).

60

4.6 Full Poisson-Boltzmann model: differential capacitance as a function of the potential for various values of the spread parameter  $\Delta$  in the planar

interface function, using the numerical solver. The distance parameter  $d$  has been set so that the interface lies 6.568 Bohr away from the surface. The full numerical PB model has been used for all calculations, using  $c_0 = 0.01$  M. 62

4.7 Same as Figure 4.5, but for the full-numerical PB model (circle and solid lines) and linearized PB model (triangles and dashed lines). The interface functions employed are the following: the planar interface (top,  $d = d_0/2 + 6.568$  Bohr, where  $d_0$  is the slab thickness, and  $\Delta = 0.470$  Bohr); the soft-sphere interface (middle,  $r = 6.568$  Bohr,  $\Delta = 0.470$  Bohr); and the SCCS interface (bottom,  $\rho_{max} = 10^{-4}$  a.u. and  $\rho_{min} = 10^{-5}$  a.u.). The top subplot also includes results obtained with the Helmholtz model (black squares and dotted line) as a comparison. The same  $d$  and  $\Delta$  parameters used for the planar interface have been employed to set the Gaussian countercharge density. 64

4.8 The differential capacitance is plotted as a function of the potential for various values of the  $c_{max}$  parameter in the MPB model. The planar interface function has been used for all calculations, with  $d = 6.568$  Bohr and  $\Delta = 0.470$  Bohr. 66

4.9 The differential capacitance is plotted as a function of the potential. Experimental data from Ref.[102] are plotted as dashed lines. Results of MPB simulations using the SCCS cavity with the original parameterization from Ref. [3] are plotted as solid lines. The value of  $c_{max}$  is set to 20 M. Red is for  $c_0 = 0.1$  M and blue is for  $c_0 = 0.04$  M. 67

4.10 The differential capacitance plotted as a function of the potential. All data refer to MPB simulations with the SCCS interface. The original parameterization from Ref. [3] has been employed, and  $c_0 = 0.1$  M and  $c_{max} = 20$  M. In the top panel, a repulsive potential between solute and electrolyte is introduced. Different colors correspond to different values of

- the  $d$  parameter, as indicated ( $w$  is set to  $0.25\text{\AA}$ ). In the bottom panel, a Morse-potential interaction between solute and anions is employed instead, with  $d = d_0/2 + 1.5\text{\AA}$  and  $w = 0.5\text{\AA}$ . Different colors correspond to different values of the  $E_{ads}$  parameter, as indicated. 69
- 4.11 The differential capacitance is plotted as a function of the potential for  $c_0 = 0.1\text{ M}$ . Experimental data[102] are shown as dashed lines. All theoretical data refer to MPB simulations with the SCCS interface function. In the top panel, results from the original SCCS parameterization [3] ('neutral'), are compared to results from the cation- and anion-specific parameterizations[27] and to results from the parameter fit to the Pt(111) PZC[47]. The value of  $c_{max}$  is set to  $20\text{ M}$ . In the bottom panel, the SCCS parameterization from Ref.[47] is employed, and the value of  $c_{max}$  varied from  $20\text{ M}$  to  $2\text{ M}$ , as indicated. 71
- 4.12 Same as Figure 4.9, but the SSCS interface has been employed in the simulations. In top and bottom panel the value of  $c_{max}$  is set to  $20\text{ M}$  and  $2\text{ M}$ , respectively. 73
- 4.13 Same as Figure 4.12, but the DC curves obtained with the MPB model and  $c_{max} = 20\text{ M}$  are compared to analogous curves obtained with the linearized PB model. 75
- 5.1 The field aware function with parameters set to demonstrate the asymmetry with a non-zero field factor. This smoothly varying function can be tuned to ignore atoms with insufficient flux, which can improve the stability of the self-consistent electrostatic calculation. 88
- 5.2 Comparison of electric flux values for each individual atom for the six molecules. A color map proportional to the absolute value of the flux is adopted. Atoms showing the largest absolute electric fluxes are identified in Figure 5.3. 91
- 5.3 Set of molecules tested. The atom possessing the maximum absolute

	electric flux is highlighted by a red circle, and its flux value is reported	92
5.4	Absolute solvation free energy error for propanoic acid in kcal/mol as a function of the mean flux threshold and the asymmetry parameter. The flux spread is fixed at 0.4, the field awareness is fixed at 0.05, while an homogenous scaling parameter $\alpha_\xi = 1.20$ is considered	94
5.5	Propanoic Acid. Unscaled interface (left) represents the spherical functions as calculated by standard SSCS. Scaled interface (right) is based off an optimized set of parameters, as reported in Table 5.1.	95
5.6	Absolute solvation free energy error for diethyl ether in kcal/mol, as a function of mean flux threshold and field awareness parameter. The flux spread is fixed at 0.6, and an homogenous scaling factor of $\alpha_\xi = 1.20$ is adopted. An asymmetry $\kappa = 1$ is selected, as in this molecule there are no atoms with negative fluxes.	96
5.7	Diethyl Ether. Unscaled interface (visualized on the left half of the molecule) represents the spherical functions as calculated by standard SSCS. Scaled interface (represented on the right half of the molecule) is based off an optimized set of parameters, as reported in Table 5.1. The effects of the field-aware scaling are small, mostly noticeable for the acidic hydrogen atom bonded to the oxygen.	97
5.8	Absolute solvation free energy error for benzyl alcohol in kcal/mol as a function of homogenous scaling $\alpha_\xi$ and field awareness parameter. The flux average and spread parameters are set to values of 2.0 and 0.5, respectively, so that only a couple of atoms are affected by the field-aware procedure.	98
5.9	Benzyl Alcohol. Unscaled interface (left) represents the spherical functions as calculated by standard SSCS. Scaled interface (right) is based off an optimized set of parameters as reported in Table 5.1.	99
5.10	Absolute solvation free energy errors of neutral water (central panel),	

	hydronium ion (top panel), and hydroxyl anion (bottom panel) in kcal/mol as functions of field awareness parameter and homogenous scaling parameter $\alpha_\xi$ . The asymmetry parameter is set to zero, +1, and -1 for the three systems, respectively. Mean and spread values of the flux thresholds are set to 1.5 and 0.6.	100
7.1	Visualization of Environ surface descriptor. Atomic centered spherical probe shown in green overlaps boundary defined by Environ soft-sphere interface (shown as a transparent 'bubble' surrounding the protein molecule). Overlap shown in bright green.	118
7.2	Free energy of cavitation of a water-like cavity immersed in a solution of water atoms. Data obtained as in section 7.3.1, with a production time of 1ns and a reduced van der Waals cutoff for simulations with smaller cells, or less water molecules.	124
7.3	Free energy of cavitation reported as a difference with respect to the final energy. Plot represents the averaged cavitation energy of seven small molecular cavities taken with respect to the cavitation energy calculation at a snapshot after 2ns of simulation time. Shaded blue region shows a standard deviation of the free energy fluctuations that occur during simulations, aggregated as described in section B.	125
7.4	Free energy of cavitation reported as a different with respect to the final energy as calculated with 20 lambda points. Plots represent the averaged cavitation energy of seven small molecular cavities taken with respect to the cavitation energy calculation at a snapshot after 2ns of simulation time. This figure demonstrates how the clustering of the trendlines (convergence with respect to total time available for sampling) is smaller than the convergence over lambda points.	126
7.5	Work done to create a cavity diameter $\sigma_2$ in TIP4P water in yellow, against hard-sphere equivalent 7.11 in blue ( $y = 0.371$ ). Pressure chosen	

- to match simulations,  $T = 300K$ . Red line denotes quadratic term in scaled particle theory to demonstrate how the contributions come into play. 127
- 7.6 The linear relationship between cavitation energy and surface area. Surface area generated by quantum surface calculation with soft-sphere interface. Red trend line generated by ordinary least squares regression with a combined training set of small molecules and proteins with 3-fold cross validation. Green trend line generated by ordinary least squares regression with the small molecules used for training data. 131
- 7.7 Best performing set of descriptors, Environ surface + volume, with lasso regression. Predicted cavitation energy plotted against simulated cavitation energy on a log-log graph. 3-fold cross validation used, graphic shows the final randomized testing set. MAE for small molecules 0.8 kcal/mol, MAE for proteins 10.2 kcal/mol. 134

## CHAPTER 1

### INTRODUCTION

#### 1.1. Materials Discovery

A vitally important role that science provides is discovery. Historically discovery has often preceded theory. For example when understanding how to exploit the world around us, humans would experiment with different substances, heating them, combining them into new potentially useful things. Once sufficient discovery had been made, scientific theory could be established, and by doing so, gaps would appear in what we know exists, predictions for us to fulfill. Once the periodic table was established in 1869 and the field of chemistry took off, discovery moved to trying to build a specific previously discovered compound in a way that could be industrialized, driven by human demand. An example of this is the development of the Haber process, in which nitrogen and hydrogen combined to form ammonia. Theory additionally provided predictions about what kinds of elements and compounds should exist, driven by the patterns that different species of atoms exhibited. Trying to find a product with specific desirable properties was much more challenging, and required a great deal of trial and error, driven by the knowledge of chemistry and physics at the time.

There are in general, two fundamental approaches to a scientific problem. One can start from first principles, from known theory which at this point in time is mostly well established, or from noticing relationships between a system with full details and one with less degrees of freedom. For example, the time taken for an object to fall from some height to the ground can be measured by considering the entire connected universe, but understanding that most of the contributions are relatively negligible, one can simply consider the object of interest, the earth and the force of gravity. It is even possible to infer the relationship of the object properties (in this case, the height above the ground when the timer starts) with the property we want to measure (the time taken) by performing experiments and finding patterns in the results.

Realistically a combination of the two should be used. This results in a number of

different philosophies for approaching problems. We desire accuracy, which implies that *ab initio* is the *modus operandi*, and approximations should be avoided whenever possible. However, we are limited by computational resources. The majority of problems of interest are simply too complex to calculate without some level of simplification.

The establishment of quantum theory in the early 20th century provided for the first time a complete exact model in which to understand the properties of arbitrary chemical combinations, however, the work required to solve the mathematics was unfathomable. The *n*-body problem is a prevalent issue that arises when attempting to apply the laws of physics to solve even simple systems. It crops up in solving planetary paths due to gravitational forces, or in solving coupled classical mechanics problems. For large enough time scales, nonlinear effects become appreciable and the time evolution becomes what is known as chaotic. It is apparent that only 2 body problems should be solvable in classical physics by considering the degrees of freedom and the conserved quantities that provide us with the ability to reduce the degrees of freedom.

With  $n$  bodies, we have  $6n$  degrees of freedom, 3 spacial coordinates, and 3 momentum components for each body. Since we only care about the motions of the bodies relative to each other, we can factor out the position and velocity of the center of mass (or indeed any reference point that moves with the center of mass of the system). This accounts for 9 of the degrees of freedom, 3 for the location, 3 for the momentum, and 3 for the rotation. Additionally we can say that the total energy of a system is conserved. For 2 bodies, we are left with 2 degrees of freedom, enough to determine a relationship between two variables. For 3 bodies, we have 8 degrees of freedom, and therefore no exact solution should be possible. This limitation in our ability to solve the physics of systems shows up again in quantum mechanics when considering the helium atom, a system with two electrons and one nucleus. To overcome this, we are forced to resort to tricks and approximations that simplify the problem. The validity of such methods depend on how well the assumptions are held. One famous approximation in quantum chemistry and a requirement for scalable calculations is the Born-Oppenheimer approximation. This starts with the assumption that the mass



of the nucleus is much greater than the mass of the electron. This greatly simplifies the wavefunction by ignoring cross terms and instead considers the product of the electronic and nuclear wavefunctions.

Two advances contributed substantially to solving multi-body quantum problems feasible. One is the rise of the digital computer, and the other is the development of Kohn-Sham equations. Kohn-Sham equations resolve the intractable nature of density functional theory by proposing that the ground state electrostatic density can be replaced by a fictitious system of non-interacting particles which can be solved numerically. This new system is exact in principle, but relies on a correct exchange correlation potential term that is, in practice, computationally infeasible to compute and hence left as an approximation. These equations were introduced in 1965 and numerical methods for computing the ground state of some quantum system did not significantly improve in the decades since. Rather, computational power that has grown exponentially even up to the present day. In fact, the entire field of computational science, from classical mechanics simulations to Monte Carlo methods, has benefited from this growth and these models continue to be used for cutting edge discovery to this day.

The advancement of computational discovery is such that now we exist in a world where theory completely precedes any form of discovery, where finding the properties of materials via trial and error can happen without any physical manifestation. The future looks even more bright. Due to the vast quantity of data we have collected over the years of performing condensed matter and chemistry simulations, we can harness the modern power of computers to synthesize mappings between known physical parameters and unknown material properties via machine learning models. Once a mapping is established, it may even be possible for the processes to be reversed. Models created using machine learning can find the structures of materials that would best fit a set of criteria.

Our understanding of materials and chemistry would not be where it is today without the many diverging paths of different research groups over the decades. Simply put, there is no best way of applying the physics that we know to an arbitrary atomic system. In

principle, if time were not a concern, humanity has reached a point where it is possible for calculate the properties of an atomic system exactly, from the properties that we know such as number of electrons, protons, and neutrons for each atom in the system. Despite the n-body problem, we can employ iterative schemes designed get closer to the exact solution over time. Yet still, if one were to consider every actor at play, even single atomic systems become intractable, and approximations and simplifications have to be made in order to make progress.

As we approach systems of larger and larger scales, greater approximations must be made, but the result is not necessarily a model without accuracy. By reducing the scope of a problem and reducing the amount of information deemed useful in learning about a system, useful simulations can be designed at any scale. For example, molecular dynamics ignores quantum effects, but models many biological systems faithfully. Models of water can be purely classical by design and yet accurately capture many of the phase transitions that exist.

## 1.2. Simulation Models

Density functional theory has for many decades been the most successful solution for dealing with periodic systems in vacuum, that is, infinitely repeating patterns of atoms. Since there are many many orders of magnitude between the atomic scale and the macroscopic scale in which we inhabit with the materials that we use, deriving properties from such assumptions as infinitely sized, or without regards to defects or edges, is quite reasonable. Density functional theory is in principle, an exact solution to the Schrödinger equation. However, to find the ground state of a system in reasonable time, one must make approximations. Even still, density functional theory, or DFT for short, does not perform well for large systems, with an approximate scaling of  $n^3$ , where  $n$  is the number of atoms. What is large you may ask? Well, even at 30 atoms, simulations slow down considerably. DFT, as it was initially conceived, is somewhat limited in the types of systems it can compute, and yet because of the periodic nature of materials, it remains to this day an incredibly successful model for discovery.

For computers to provide insights into chemical interactions, an understanding of larger scale systems is necessary. A common fascination is an isolated molecule surrounded by water molecules. Away from areas of interest, it is necessary to simulate many water molecules and since they typically exist as some kind of medium, their configuration is constantly changing, which adds a lot of complexity to the problem. The solution adopted by much of the computational chemistry community was to ignore the quantum effects whenever we did not need detailed information about the interaction between electrons and instead focus on a classical model. In its simplest form, a classical molecular dynamics simulation consists of a collection of atoms, each with pairwise interactions that serve to approximate every electrostatic effect. Like in DFT, assuming periodicity resolves many of the scaling issues that simulations might have and allows us to make macroscopic conclusions from systems with thousands of atoms.

Classical models certainly have their advantage of their purely quantum counterparts. They scale much better with system size, and provide computational scientists with the ability to gain insight into systems of higher complexity. Much of the information is lost, however, and in practice, both tools find themselves with their own distinct applications.

Further approximations can be made to reach areas where classical methods fall short. Take a single species liquid for example, like water. In order for Molecular Dynamics, or MD, to derive the effects of water on larger systems like a dissolved material or the surface of a material, many water molecules need to be considered, and as stated before, many snapshots need to be considered, because the configuration of a liquid is continuously varying with time. In order to derive any time based property, one needs sufficient snapshots for statistical accuracy. However it is clear that if were to statistically average the configurations of the water molecules, they would exist continuously over all space. This effect breaks down in heterogeneous regions, but in a large medium of water, one can fairly confidently represent the system without atoms entirely. The substitute is a field, or a continuum, where the properties are permitted to vary continuously over space. Such an approximation finds a lot of application in finding macroscopic properties of materials and media like elasticity and

plasticity.

Multi-scale models serve to bridge the gap between these models. The general idea is that one should be able to split up an arbitrary system into subsystems separated by boundaries. If we are interested in only a single subsystem, the rest of the system can be represented by a more approximate model, and therefore one effectively gains the best of both worlds. A detailed representation of the area of interest and the ability to simulate a larger system without encountering computational scaling issues. In this writing, I will be focusing on the combination of DFT and continuum models.

Continuum solvation models, or implicit solvation, are multi-scale models that combine DFT and continuum mechanics in order to produce fast simulations of systems that involve an interface with some medium. Since their adoption, the continuum idea has been expanded to include the electrochemical double layer that exists due to the presence of ions in solution.

In the field of continuum solvation and multi-scale models exist a range of implementations that take slightly different approaches. The work of Andreussi and others currently exists in a software package called Environ, and is in continual development primarily by members of the Materialab group at the University of North Texas, led by Dr. Andreussi himself. Environ works on top of Quantum ESPRESSO, the most widely used open source DFT software suite in active development. Environ contains continuum solvation models that couple with the DFT drivers of Quantum ESPRESSO. The models themselves are robust and contain very few artificial parameters.

One of the most important distinctions in a model is in the definition of the interface that separates the vacuum DFT region with the continuum represented implicit solvent. This can be based on the density of the electrons around a solute, the atomic positions in the solute, or a combination of the two. In general, there is a tricky balance between adding features to a model and thus increasing its transferability and accuracy, and adding too many features, such that there are unnecessary redundancies, and extrapolation works very poorly.

### 1.3. Continuum Solvation Specifically

Let's take a step back and consider the importance of simulations with interfaces. Humanity exists as the self proclaimed dominant species of Earth, and thanks to an ever growing shared knowledge-base called science, we have learned to outpace any natural population inhibitions, and effect our environment on a planetary scale. As a collective however, we are arguably merely children that for too long have lacked the wisdom to wield the power that we possess in a way that does not compromise our own continued evolutionary success, and without concern to the diversity of life that exists around us.

The general trajectory is one towards automation, making survival more and more trivial by driving up the efficiency of processes, the amount of actions possible by a single human in a given span of time, at the cost of energy. This began with the discovery that biomass was combustible and therefore a source of energy, and the domestication of animals, and moved progressively to more dense forms of combustible fuels, like coal and oil.

Our current situation, as of 2021, is somewhat dire. Atmospheric carbon dioxide, the gas that accounts for the largest quantity of emissions from combustion, has increased from 260–270ppm before industrialization [105] to 420ppm today [92]. The result is a global change in air and surface temperatures [43, 61], and trends are accelerating. A somewhat simplistic way to look at the problem is to focus solely on these CO<sub>2</sub> emissions, the rate in which they are increasing annually, and the current transition away from combustible fuels. Annual carbon dioxide emissions have increased by 45% in the last 20 years whereas the global energy consumption increased by 42% [83, 84]. The increase in adoption of renewable energy sources, whilst encouraging, pales in comparison to the usage of coal, oil, and gas, which combined, continue to see increased usage year after year. The driving factor behind the rise in renewable popularity has been their economic viability, which is in turn driven by materials innovation. Therefore not only is materials discovery useful for the advancement of all applications that we use today, but it is critically important in the fight to mitigate the effects of climate change, a task for which our time becomes increasingly constrained.

There are many applications for continuum solvation models, and the situation in the

world right now is driving current developments very much towards the research of materials used in renewable technologies. A significant challenge faced in chemistry is the search for catalysts. When posed as a computational problem, continuum solvation fits remarkably well [103, 104]. We have a system composed of a catalyst, usually existing in some liquid medium that is capable of facilitating some chemical reaction. This opens up more options for what reaction pathways are viable, and provides economically viable or perhaps even more sustainable alternatives to existing pathways.

One potential fuel source leverages the mechanism of photosynthesis [19, 18]. Another uses electrolysis; fuel cells exist as a potential answer to our current dependence on combustible fuels. Hydrogen fuelled and battery powered electric vehicles are the front-runners in the race to replace the gasoline powered vehicle. Each has its advantages and disadvantages, the latter of which can be mitigated by funding and research. As of 2021, hydrogen fuel cells have found a niche as an alternative for larger vehicles, where the capacity limitations of batteries become relevant.

A major barrier preventing more widespread adoption of these sustainable replacements is the economic viability of hydrogen. Currently the fuel cell relies on rare metals to catalyse the electrochemical reactions of water and its constituents, of which Platinum has been shown to be the best. A so-called 'volcano' plot (a graphical representation of the Sabatier principle) was initially used to illustrate how catalyst choice is typically a compromise between being too strong and therefore preventing dissociation of the product, and too weak such that the adsorbent does not bind to the surface. A similar trade-off can be seen in the materials used for the electrocatalysis of water. Take the hydrogen evolution reaction for example [98]. By studying the bond strength of the adsorbed hydrogen vs. the exchange current for the reaction, one sees the same apparent volcano plot shape, with platinum at the top, with the minimum exchange current, and a bond strength that is not too strong and not too weak.

2D materials are an exciting new development in materials discovery [68, 53, 54] that have different properties to their bulk counterparts, of which their use in electrocatalysis is of

particular interest. Continuum solvation models are capable of assessing the stability of 2D materials in water (or indeed any solvent). They can also test for the stability of adsorbates that bind to 2D catalytic surfaces. A series of papers by Naiwrit Karmodak explores the possibilities of these materials as a replacement for the current platinum catalysts that lack economic viability.

The study of thermodynamic properties of electrodes in solution is somewhat involved and requires a few pieces on top of the standard continuum solvation model. The electrochemical double layer is represented by a continuum [74], and the charge and adsorbate coverage is varied as part of a grand canonical approach.

Continuum solvation has also been used to investigate the dependence of permittivity on catalytic activity, in order to find new directions to take to solve the hydrogen fuel cell problem [32] and to investigate the properties of Pt nanoparticles instead of its bulk equivalent for economic viability [87]. Montemore [67] applies solvation to investigate the adsorption of alkyls on a copper catalyst. These intermediates exist as part of the CO<sub>2</sub> electroreduction and the aqueous Fischer-Tropsch process (converting fuels from natural gas and other forms of biomass).

## CHAPTER 2

### CONTINUUM SOLVATION

A solvated system consists of a solvent and some dissolved molecules. The solvent is typically some liquid medium that consists of a single type of molecule, for example, water is a polar solvent. Water can dissolve polar molecules due to their ability to connect to the hydrogen bonding network that the polar water molecules typically form. This electrostatic interaction keeps the energy of the solvated system below the energy of the separated system. The free energy of solvation (or the energy required to dissolve a solute in some solvent) is given as the sum of many terms:

- Electrostatic - all electrostatic terms are solved by the Poisson equation, once solved, the energy can be computed.
- Cavitation - the energy required to create a vacuum filled cavity of some shape inside the solvent medium.
- Repulsion - the continuum counterpart of the short-range interactions induced by the Pauli exclusion principle.
- Dispersion - the van der Waals interactions.
- Vibrational/Rotational changes
- Volume changes in the solute Hamiltonian

Continuum solvation models must solve a classical electrostatic problem in order to account for the field representation of the solvent. The region of space that is populated by solvent molecules has some constant static permittivity greater than one, and the region populated by the solute is treated as a vacuum, with a static permittivity of one. Since the function for static permittivity is not constant over all space, the electrostatic problem is not so simple to solve for,

$$(2.1) \quad -\nabla \cdot (\epsilon(\mathbf{r})\nabla V(\mathbf{r})) = 4\pi\rho_M(\mathbf{r})$$



where  $\rho_M(\mathbf{r})$  is the solute charge distribution and  $\epsilon(\mathbf{r})$  is the permittivity. Hence, for a given potential  $V$ , one can solve for the charge distribution of the solute due to some solvent.

### 2.1. Polarized Continuum Model

One of the most commonly used methods in computational chemistry for the modeling of solvation effects is the polarizable continuum model [65, 66], PCM for short. PCM splits a system into the solvent medium and the solute (often a single dissolved molecule). The boundary between solute and solvent is a hard interface built up from interlocking spheres centered at each of the atomic positions of the solute. The radii of these spheres is set by the van der Waals radii. The cavitation energy is approximated as the surface area of the interface. The electrostatic problem is approached by considering the effect of the solvent through the interface. Computationally, it is possible to break down the interface into many small tesserae, each of which has some surface charge. The potential from eq. 2.1 can be written as the sum of two terms, the potential due to the solute and a potential term due to the surface of the interface.

To solve for the solute region, which is a quantum mechanical problem, one might have an effective Hamiltonian for the solute that consists of two terms, the Hamiltonian for the solute in vacuum and a new term representing the electrostatic interaction between the solute and the solvent.

$$(2.2) \quad \hat{H}^{\text{eff}} = \hat{H}^0 + \hat{V}^R$$

Continuum solvation models present themselves as a viable approximation to the explicit solvation achieved by Molecular Dynamics, which simply becomes too expensive in bulk. Pure DFT on the other hand, does not describe solvents well due to limitations in the functionals.

From the other side of the scientific community, the computational condensed matter physicists had embraced DFT for quite some time,

Arguing for a particular continuum model from a design standpoint is somewhat out of scope for this dissertation, due to the required theory, and therefore I shall simply consider the current state of continuum solvation models, how well they perform, and what can be done to improve a couple of the more recent models that have seen success in the literature. Much of my work follows the work of Andreussi, Fisicaro, and their collaborators.

## 2.2. Self Consistent Continuum Solvation Model

The Self Consistent Continuum Solvation model [3], or SCCS, is a reformulated version of the solvation model of Fattebert and Gygi [29]. The original motivation for Fattebert and Gygi in 2002 was to incorporate solvation into DFT, by modifying the original Kohn-Sham potential into a form that accounted for the additional effects of the solvent on the solute. In the original Kohn-Sham formulation, the energy functional consists of four terms, the kinetic energy, the energy due to an external potential, the Hartree energy that accounts for the electrostatic effects, and the exchange-correlation energy.

To account for solvation, the Hartree energy functional becomes

$$(2.3) \quad \frac{1}{2} \int \rho\phi[\rho]d\mathbf{r}$$

where  $\phi[p]$ , the electrostatic potential, can be solved from the Poisson equation, eq. 2.1. The modified Kohn-Sham equations are solved for in all space inhabited by the solute and the solvent, and the effect of the solvent is captured by the permittivity functional that can vary over space. Unlike PCM, the shape of the permittivity functional cannot be discontinuous, because of the way plane wave DFT is typically computed, and therefore one considers a smoothly varying function instead. Fattebert and Gygi go one step further and redefine the solute region by the electronic density rather than the atomic positions, the rationale being that this electronic function is a better representation of how the solute interacts with the outside world. The electrostatic free energy can be written as,

$$(2.4) \quad F[\rho^{\text{el}}, \{\mathbf{R}_i\}] = \int \left( \rho^{\text{el}}(\mathbf{r}) + \sum_j z_j \delta(|\mathbf{r} - \mathbf{R}_j|) \right) \phi(\mathbf{r}) d\mathbf{r} - \int \frac{1}{8\pi} \epsilon(\mathbf{r}) |\nabla \phi(\mathbf{r})|^2 d\mathbf{r}$$

where  $\phi^{\text{el}}$  is the electrostatic density due to the electrons,  $z_i$  are the ionic charges, and  $\epsilon(\mathbf{r})$  is the dielectric permittivity of the system. The generalized Poisson equation is recovered by taking the functional derivative of the free energy expression with respect to the electrostatic potential. The choice of the permittivity functional is somewhat arbitrary, needing to satisfy only a few constraints, and needing to be stable in applications. The functional proposed by Fattebert and Gygi was a simple switching function that relates back to,

$$(2.5) \quad \epsilon_r(\mathbf{r}) = s(\mathbf{r})(\epsilon_0 - \epsilon_b) + \epsilon_b$$

where  $\epsilon_b$  is the permittivity of the bulk solvent,  $\epsilon_r$  is the permittivity at an arbitrary point due to the solvent and  $\epsilon_0$  is the permittivity in vacuum.  $s(\mathbf{r})$  is a rescaling of the permittivity at an arbitrary point to fit some standard mathematical switching function that ranges from 0 to 1.

$$(2.6) \quad s(\rho(\mathbf{r})) = 1 - \frac{1}{2} \left( 1 + \frac{1 - (\rho(\mathbf{r})/\rho_0)^{2\beta}}{1 + (\rho(\mathbf{r})/\rho_0)^{2\beta}} \right)$$

with  $\rho_0$  describing the position of the solvent-solute interface, and  $\beta$  determining the smoothness of the function.

With the switching function defined, one can relate the electrostatic free energy to the Kohn-Sham equations. This unified framework determines the Kohn-Sham potential and the forces,

$$(2.7) \quad V(\mathbf{r}) = \int \frac{\delta F[\rho^{\text{el}}]}{\delta \rho^{\text{el}}(\mathbf{r})} d\mathbf{r}' = \int \frac{\delta F[s]}{\delta s(\mathbf{r}')} \frac{\delta s(\mathbf{r}')}{\delta \rho^{\text{el}}(\mathbf{r})} d\mathbf{r}'$$

$$(2.8) \quad \mathbf{f} = \int \frac{\delta F[\mathbf{R}_i]}{\delta \mathbf{R}_i} d\mathbf{r} = \int \frac{\delta F[s]}{\delta s(\mathbf{r})} \frac{\partial s(\mathbf{r})}{\partial \mathbf{R}_i} d\mathbf{r}$$

The original model for the first time brought the already well developed DFT methods from the condensed matter community and the continuum solvation methods originally pushed by the quantum chemistry community together. Since then, interest in this new family of methods has grown exponentially. Naturally the next step would be to consider the non-electrostatic interactions. In 2006, Scherlis, Cococcioni, and Marzari at the Massachusetts Institute of Technology collaborated with Fatterbert and Gygi, now at Lawrence Livermore National Laboratory to add a cavitation energy term, defined as the work involved in creating a cavity inside the solution whilst ignoring any of the solvent-solute interactions. Cavitation is of particular interest and will be revisited later in an attempt to improve the current state of continuum solvation models. At this point, a number of prevalent models were considered. What was chosen was a reasonable estimation,

$$(2.9) \quad \Delta G_{\text{cav}} = \gamma S(\rho_0)$$

where  $S(\rho_0)$  is the surface of the interface defined in solving for the electrostatic energy contribution, and  $\gamma$  is the macroscopic surface tension of the solvent. This formulation is remarkably simple and effective. The original formulation of Fatterbert and Gygi and the possibility of incorporating their methods into DFT was hindered by the requirement of a fast multigrid solver (an algorithm that solves for the Poisson equation via a hierarchy of discretizations) and robust choice for the permittivity functional.

This work, published in 2012 by Andreussi, Dabo, and Marzari begins by recasting the problem in terms of induced polarization charges, suggests an improved formulation of the permittivity functional, provides an improved algorithm for solving the Poisson equation, and improves on the non-electrostatic energy contributions. Rather than a switching function based off the density directly, the logarithm of the density was considered, since the density due to electrons decays exponentially when moving from the interface to the

solvent. Alternative switching functions were tested and a trigonometric one was chosen as the most stable option due to the reasonable nature of its gradient function. The logarithm of the electronic density was chosen in order to account for the exponential behaviour of the density away from the solute.

$$(2.10) \quad s(\rho(\mathbf{r})) = 1 - \frac{1}{2\pi} \left[ 2\pi \left( \frac{\ln(\rho_V) - \ln(\rho)}{\ln(\rho_V) - \ln(\rho_\Lambda)} \right) - \sin \left( 2\pi \left( \frac{\ln(\rho_V) - \ln(\rho)}{\ln(\rho_V) - \ln(\rho_\Lambda)} \right) \right) \right]$$

where  $\rho_V$  and  $\rho_\Lambda$  are the maximum and minimum bounds for the electronic density in the switching function, respectively. These parameters determine where the interface sits, and should therefore change depending on the interaction between the solute and the solvent molecules. These interactions cannot be captured well with an implicit solvent model, and therefore any local determination of the interface function (that is, a function that doesn't take into account its surroundings) will suffer from some minor shortcomings. Minor reparameterizations will be needed in order to account for changes in the system, like the charge on the solute, any electrostatic interactions between the solute and a polar solvent, or the structuring of solvent molecules around a solvent.

The expression for the interface is continuous in its first two derivatives, which is useful for solving analytic terms in the Poisson equation. In order to parameterize an interface model, one can calculate solvation energy for molecules with known experimental values, and fixed charge and tune the parameters in order to minimize the error between the experimental values and the values calculated by the simulations.

The model is self-consistent due to the choice of an iterative algorithm in solving the Poisson equation. In addition, the Poisson equation as stated in eq:2.1 is recast into an expression in terms of the polarization density. This is in the same vein as the PCM

There are other options for solving for the electrostatic problem, each with its own advantages and disadvantages

- FFTs
- Multigrid

- Wavelets

### 2.3. Soft-sphere continuum solvation

As we've seen so far, there are two general pictures for determining the shape of the interface around a solute. PCM takes an atomic centered approach, and SCCS takes an electronic density approach. The soft-sphere model aims to take the atomic centered approach of PCM and apply it to the combined implicit solvent / DFT approaches that began with Fattbert and Gygi. The framework is therefore similar to the SCCS, with an interface function

$$(2.11) \quad s(\mathbf{r}, \mathbf{R}_i) = 1 - \left[ \prod_i h_i(\{\xi\}; \|\mathbf{r} - \mathbf{R}_i\|) \right]$$

where  $\{\xi\}$  is a set of parameters describing the spheres. Visually, this looks almost identical to PCM, where we have interlocking hard-spheres, however, as the name suggests, these are soft-spheres represented in three dimensional space. Since the derivative of the interface with respect to the atomic positions is necessary to solve the Kohn-Sham equations, the functions must be differentiable, and so the interface is instead a product of interlocking radial error functions.

These error functions,  $h$  can be written as,

$$(2.12) \quad h_i(\alpha_\xi, R_i^{\text{vdW}}, \Delta_\xi; \|\mathbf{r} - \mathbf{R}_i\|) = \frac{1}{2} \left[ 1 + \text{erf} \left( \frac{\|\mathbf{r} - \mathbf{R}_i\| - \alpha_\xi R_i^{\text{vdW}}}{\Delta_\xi} \right) \right]$$

which has two tunable parameters (and one 'physically defined' parameter),  $\alpha_\xi$ , a global multiplicative scaling on the size of the spheres, which couples with a set of radii (typically van der Waals radii,  $R_i^{\text{vdW}}$ , so each atomic species will have a different 'interaction' radius), and a spread function  $\Delta_\xi$ , which controls the region of polarization charge.

## 2.4. Model parameterization

By design these models work on neutral molecules. Since they are built on DFT, there is a molecular size limit for a reasonably timed simulation. Performance-wise, the SCCS and the soft-sphere models have comparable simulation times to vacuum simulations. Therefore, a small set of organic molecules are chosen (240 total) such that molecules with properties that are not captured by the implicit model do not skew the distribution of solute-solvent mean interaction distance significantly.

Extrapolating away from the parameterization set is somewhat tricky. Previously discussed was the effect of solute charge on the simulations. Consider a polar solvent like water. If a solute is positively charged, water molecules close to the solute will reorient themselves such that the positively polarized part will face away from the charge and the negatively polarized part will face towards the charge. There is also an electrostatic attraction between the solvent and the solute. This combined effect means that the mean distance between solvent and solute decreases, and in the reverse scenario, where the solute is negatively charged, the effect is a lot more pronounced.

The mean absolute error in aqueous solvation free energy, as compared to experimental values, is a popular test to determine how well various competing interface models perform. The SCCS model for instance, performs well for neutral solutes, with an error close to chemical accuracy of 1 kcal/mol. However, further analysis shows that this error is highly dependent on the functional groups possessed by certain molecules, as expected from the above discussion. In particular, mean errors for alkenes, ketones, and aldehydes are shown to be well below 1 kcal/mol in many cases, whereas amines, ethers, and acids show errors of up to 4-5 kcal/mol. When applied to cations and anions, this same model has errors of around 2 kcal/mol for cations and 17 kcal/mol for anions, an order of magnitude higher. In contrast, the SSCS model performs comparably for neutral molecules, and again worse for cations (6 kcal/mol) and anions (10 kcal/mol).

In general, a set of cations and anions is harder to fit than neutral molecules, since there is more variance in the interactions that don't get captured by implicit solvent. A

simple reparameterization of the SCCS model is possible by varying the parameters for a training set consisting of small cation or anion molecules, and minimizing the mean absolute solvation error. These parameters are then fed into the simulation according to the a priori knowledge of the charge state of the solute. Indeed, this does bring down the mean absolute solvation error down by a small amount for cations (2.55 to 2.26 kcal/mol) and more significantly for anions (17.4 to 5.54 kcal/mol). Having higher accuracy for ions in particular greatly expands the scope of applications for continuum solvation, from capacitance measurements to reaction intermediates that consist of ions.

The effect of this is to shrink the interface, thus reducing the volume of the quantum-mechanical region. This captures the physical effect of the solvent interacting more closely to the solute molecule. However, it has to be noted that such a strategy may have limited application for simulations involving different species, in which only the global charge state of the system is assigned a priori, while the specific charge state of the each constituents is unknown and, possibly, depends on the embedding environment.

Reparameterization can also account for certain functional groups. If it is known that a particular species of atom behaves differently as part of a functional group, one might adjust the radius of this atom based on chemical intuition. This approach relies a lot on the human, and does not generalize gracefully.

## 2.5. Non-local models

These models perform well when simulating small systems, however, larger systems with more complicated shapes incur some interesting problems when trying to decide which regions ought to be solute and which ought to be solvent. Consider a folded protein, the structure is compact due to the molecule's ability to weakly self interact and minimize its own interaction with its surroundings (lower surface area). In the process, small cavities will often form within the protein due to the complexity of its shape. These cavities might be too small for solvent to fit, but the model as currently defined thinks of the solvent molecule implicitly, and so it fills in the volume not explicitly occupied by the protein anyway.

A local representation of the solvent does not work here anymore and we have to



consider the neighborhood around each atom in order to get an idea of how to place the interface more generally. By exploiting the operation of convolutions, one can identify whether the region of space around a point is mostly part of the solvent region or mostly part of the solute region. We call the new functional the filled fraction,  $f^{\text{ff}}(\mathbf{r})$ , which is defined similarly to the interface function  $s(\mathbf{r})$ ,

$$(2.13) \quad f^{\text{ff}}[s(\mathbf{r}')](\mathbf{r}) \equiv \int s(\mathbf{r}')u(\{\xi\}; \|\mathbf{r} - \mathbf{r}'\|)d\mathbf{r}' = s * u(\{\xi\}; \mathbf{r})$$

where  $a * b(\mathbf{r})$  denotes the convolution between two functions. As in the soft-sphere representation, the spherical probe  $u$  that determines the region of space around each point  $\mathbf{r}'$  is sized depending on the system, in this case however, rather than depending on the atoms in the solute, depends on the solvent atoms,

$$(2.14) \quad u(\alpha_\xi, R^{\text{solv}}, \Delta_\xi; r) = \frac{1}{2N_u} \left[ 1 + \text{erfc} \left( \frac{r - \alpha_\xi R_i^{\text{vdW}}}{\Delta_\xi} \right) \right]$$

Note the similarity between this and soft-sphere function. We use an error function to represent the spherical shape of the probe, define a smooth interface in order to produce well behaved derivatives, and add a normalization factor,  $N_u$  which can be numerically solved for by integrating the volume of the probe. This updated interface function, the solvent-aware function [4] has a solvent-aware parameter that tunes the size of the spherical probe, thus adjusting how aggressively gaps are filled in the definition of the interface.

Another use for non-locality is in the definition of the electric field, which can be used to infer atomic information about electrostatic charge. A more recent paper [100], which is itself included as part of this work describes a new interface definition that relies on the electrostatic flux, which is itself a product of the electric field and the existing soft-sphere interface function [31], which serves as a definition for the surface on which a normal component for the electric field can be computed. The result is a new field-aware model which accounts both for possible charges that might cause electrostatic interactions between

water and the system of interest, and for the asymmetry that arises due to the asymmetric nature of polarizable solvents like water.

## 2.6. Additional Energy Terms

Clearly solvents provide additional effects beyond electrostatics on the embedded system. Some of these additional terms in the free energy calculation can be represented as continuum effects. Typically they use the same interface function as the one used for electrostatics. External pressure uses the quantum volume, which can be calculated from the interface function,

$$(2.15) \quad V = \int s(\mathbf{r}) d\mathbf{r}$$

This external pressure enters an enthalpy term,

$$(2.16) \quad G^{\text{PV}} = PV$$

A cavitation energy term can be written in terms of the quantum surface, defined again in terms of the interface function,

$$(2.17) \quad S = \int \|\nabla s(\mathbf{r})\| d\mathbf{r}$$

which relates to cavitation by,

$$(2.18) \quad F^{\text{cav}} = \gamma S$$

## CHAPTER 3

### SIMULATION SOFTWARE

#### 3.1. Quantum ESPRESSO

Quantum ESPRESSO [40, 38] (QE) is an open source suite of codes for performing electronic structure calculations using DFT, plane waves, and pseudopotentials. Being open source, it also serves as a primary platform for scientists to develop new models in code. Of particular interest to the topics covered here is the plane wave (or PW) program provided by Quantum ESPRESSO, and two of the calculations that it performs. The basic calculation is a self-consistent calculation (SCF) that iteratively solves the Kohn-Sham equations for a plane wave basis set. The geometry optimization calculation (relax) additionally calculates the forces on each atom of a structure and attempts to iteratively find a stable configuration (such that the forces on each atom are close to zero).

#### 3.2. Environ

Environ [2, 6] is an open source suite of codes that follows in the same vein as Quantum ESPRESSO (having started out as a plugin to QE). It uses DFT, plane waves, and pseudopotentials to perform electronic structure calculations with environmental effects. Of interest to the topics covered here are the SCF and relax calculations. Environ relies on some kind of driver to perform these calculations. In version 1.0, the software used QE to solve the Kohn-Sham equations, and additional terms calculated from functional derivatives would pass from Environ to the QE driver every time the electrons or ions needed to be updated. In version 2.0, the software was made more general such that any driver could hand Environ information about the ions and electronic density, as well as information regarding the partitioning of the real-space and reciprocal grids.

The majority of development covered here is implemented in Environ and either already released, or expected to release in the near future. This includes an implementation of the field-aware model, faster computation for the derivatives of soft-sphere functions, and various diffuse layer models. New developments may also be implemented outside of Environ,

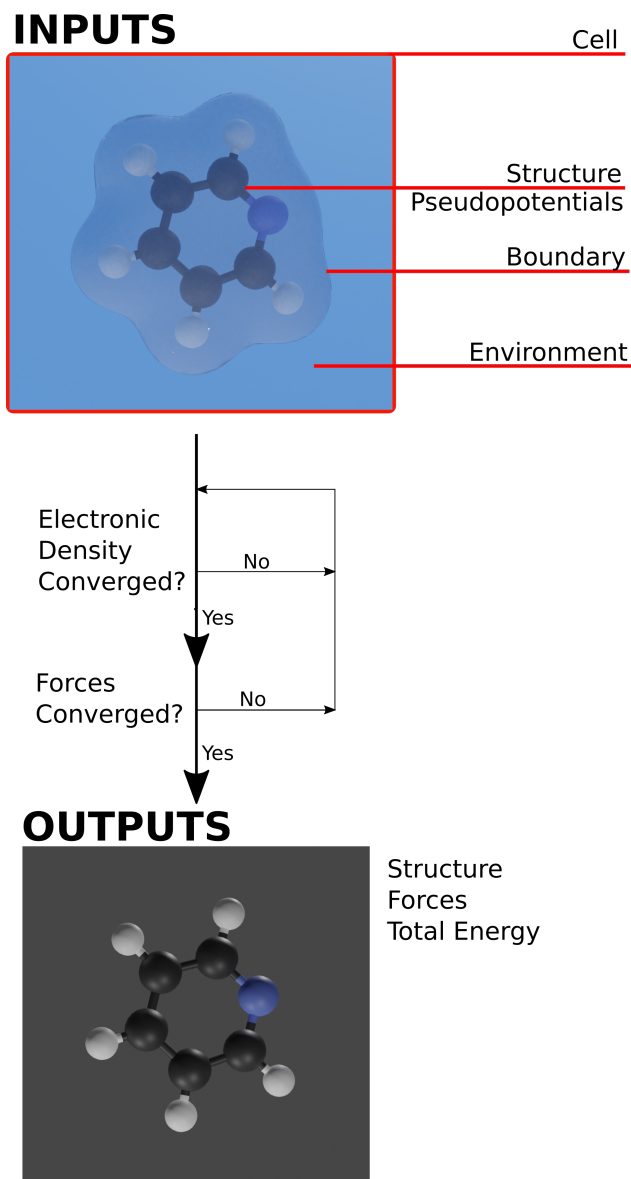


FIGURE 3.1. Workflow of a typical calculation in Environ. In order to implement new interface definitions, functional derivatives for the energy and force terms need to be derived.

either as a consequence of the pythonization efforts to be discussed later, or as automation and postprocessing tools.

As it stands, Environ is designed with modularity in mind, possessing for instance, the ability to use different methods of performing calculus on grids. Maximizing transferability does have the drawback of a need to keep the base code as general as possible and leaving

the calculations for specific applications to the user. This works well for new research, however, users may often wish to repeat calculations made previously by others but with a few parameter changes, perhaps a different set of input structures.

AiiDA (which will be discussed later) solves this by providing developers with the ability to share their workflows in a standardized plugin registry. The widespread adoption of Environ and sister codes is contingent on a user friendly environment for both development and application.

### 3.3. Pythonization

The current landscape of multi-scale models consists of various software tools developed in house. Each has advantages and disadvantages that allow them to fit a specific niche. However, we reach a point where there is not much direct improvement to do, and one must instead consider the modular nature of these models instead. Simulation of more involved systems becomes possible when combining the models that exist, such that each acts on a particular region of the simulation.

An example is the concept of combining a continuum embedding scheme with a frozen density embedding scheme. This has a lot of potential to simulate regions that can contain many water molecules (an advantage of continuum embedding) whilst paying close attention to specific interactions of the water molecules close to the quantum mechanical region (an advantage of frozen density embedding). Alone, both methods lose out on either accuracy or computational speed, but together, and many possibilities open up for applications.

The question is therefore, how does one go about combining two code bases developed by different research groups in a space where software developers are a rare commodity? Python is a three decade old interpreted language that has gained extraordinary popularity within academia due to accessibility, code readability, development speed, and interoperability with other languages, specifically Fortran and C. Where Python lacks in speed, compiled C and Fortran libraries can be used. Numerical recipes and algorithms that deal with vector and matrix algebra can all be exposed to Python, and the scripting nature of the language allows scientists to conveniently perform calculations with much easier customizability than

a program might provide. Such modularity of computational packages did exist prior to the existence of Python. However, the accessibility of Python has created a user community, that to date is unmatched in the world of computational science.

`f2py` [73] is a Python module maintained by `numpy` (as of 2021) that automatically creates a Fortran library out of code that can be imported into python and used directly through functions within the imported module. This typically works well when an already written function needs to be exposed to Python, or if a function relies on expensive vector operations that Python cannot handle effectively. `numpy` can handle many of the scenarios that exist due to the latter, and `scipy` has a large library of wrapped functions for dealing with vectors and matrices (much of the code originally being written in Fortran). The migration of high performance code is also possible, since Python has an MPI module that can interface with a `f2py` created library containing parallel subroutines. `f2py` is limited to classic Fortran, before the addition of object-oriented keywords and complex datatypes starting in Fortran 90.

Fortunately, work has been done to circumvent this restriction, existing as an unofficial `f2py` patch known as `f90wrap` [55] developed primarily by James Kermode. This tool handles the majority of additions to modern Fortran, and without much preparation, even large Fortran code bases can be transformed into compiled Python modules. There are a number of philosophies going forward. `f90wrap` can be used out of the box to generate a wrapper, but this wrapper is designed for use by a Fortran program, not by Python. In Python, we may want more control over the exact flow of the code, or over the objects that normally get autogenerated by `f90wrap`. It may therefore be beneficial in some cases to modify the original codebase in preparation for a tool like `f90wrap` or `f2py`.

For maximum freedom, a modern Fortran to classic Fortran parser can be developed

`QEpy` [88] is a pythonized Quantum Espresso maintained by Xuecheng Shao. It utilizes `f90wrap` allowing researchers access to the main Quantum Espresso drivers in Python. QE as a driver opens up much more possibilities for codes that rely on it for DFT calculations. By pythonizing Environ, one removes the limitations imposed by the design of Quantum

Espresso and the way plugins are able to work with it.

### 3.4. DL POLY

DL POLY [90] is a classical molecular dynamics simulation software that is often chosen for its extensibility. The code is free for academic use and written in Fortran. Performing cavitation energy calculations is something that works out of the box for most classical molecular dynamics codes, and therefore some amount of modification to the source code is necessary. More recently available options written in flexible languages like Python were not considered for the task, however, in the long term, these might prove to be better options.

### 3.5. AiiDA

AiiDA [49] is an open-source Python infrastructure that helps researchers automate, manage, share and reproduce complex workflows associated with modern computational science. The power of this infrastructure lies in the ability for developers to write plugins that not only integrate their software with AiiDA, but also provide automation tools for these plugins that assist with preprocessing and postprocessing, tasks that some of the older, legacy codebases struggle with.

The database serves to centralize information in a way that greatly benefits research that relies on multiple resources for computation, an increasingly common necessity. AiiDA runs as a background process that can monitor and send jobs from a local machine to a remote computing machine. It can also check for the completion of these jobs and automatically parse any results back into the database. The result is a connected graph of input and output data that can be shared for convenient collaboration and validation, and the query system helps to efficiently reference previously computed information.

The project has attracted a community of developers primarily in the field of materials discovery to write plugins for various codes. AiiDA serves as a central hub for sharing data to and from these codes that are typically maintained by separate research groups. The result is increased compatibility between codes and the development of workflows that can benefit from the combined power of different codebases.

Since Environ historically was integrated into Quantum Espresso, the development of a plugin for Quantum Espresso by the AiiDA team itself has some compatibility with Environ, however it is limited and more recent developments of Environ have pushed the necessity of a completely independent plugin for Environ.

`aiida-environ` [99] is a plugin currently released in beta status. The vision is to provide users with tools for setting up some of the more complicated simulations that are possible in Environ. Since Environ still depends on the Quantum Espresso plugin, they share similar input and output. The possible errors that occur in Environ are a superset of the errors that occur in Quantum Espresso. It is therefore reasonable to design `aiida-environ` using dependencies from the `aiida-quantum.espresso` plugin, exploiting object oriented programming ideas.

Currently implemented in `aiida-environ` is a `CalcFunction` for the PW program, that extends the `aiida-quantum.espresso` `CalcFunction` for the PW program. Environ requires an extra input file for the Environ parameters. The plugin separates the cards that would usually be included in this extra input file. Therefore, most of the parameters exist in `environ.parameters`, a Python dictionary, and the rest exists as separate inputs. For example, the external charges have their own datatype defined by the plugin, and therefore they get added separately. Note that as of the time of writing, not all the Environ inputs are exposed to the plugin.

`aiida-environ` contains a parser that for the time being is an adaptation of the `aiida-quantum.espresso` parser, due to the fact that the Environ output is written to the human readable output file of Quantum Espresso, but not to the XML file. Optionally, extra output from Environ is directed to a separate file, and the plugin is currently able to parse some, but not all of this output.

The plugin is designed to be as modular as possible. Since the focus is on sharing pre- and post-processing tools with users of Environ, automated workflows are written in Environ, and any calculations that these require are separated into helper classes and functions that do not store their calculations in the database, and `CalcFunctions` that do. A couple of






FIGURE 3.2. Solvation workflow implemented in AiiDA consisting of two PW calculations, one in vacuum and one in solution.

these workflows are adapted straight from the `aiida-quantumespresso`, that provide extra automation around handling errors in the `scf` and `relax` calculations in the PW program.

### 3.5.1. Solvation

The `solvation` workchain automates a common calculation performed by Environ, the solvation free energy. The workflow runs two Environ simulations, one in vacuum and one in solution. It then performs a simple postprocessing calculation, taking the difference between the final energy from both simulations.

The workchain expects one of two possible inputs by the user.

- (1) An `environ-parameter` dictionary as per a regular Environ calculation
- (2) An `environ-parameter` dictionary with shared variables and one/two dictionaries for custom vacuum/solution input.

### 3.5.2. Grand Canonical

The grand canonical workflow automates grand canonical simulations [47] that find some of the thermodynamic properties of electrochemical interfaces of metallic electrodes in aqueous environments. The workflow is fairly involved and consists of a number of steps. The initialization, the preparation of all the required simulations to make a calculation for a single material, the submission of the simulations and parsing, and finally the postprocessing, where the Legendre transformation takes place.

The workflow takes in a number of inputs,

- `EnvPwBaseWorkChain`. A standard Environ workchain (which should involve a basic calcfuction with some automated error handling etc., mirroring the PW base workchain that exists in `aiida-quantumespresso`). This can be mostly defined by the user directly using the builder for convenience (see the example).
- `Vacancies`. A list of coordinates for the possible vacancies. This might be too general, and we may in the future want this as a datatype (since it's used in other workflows too). General due to the fact that we assume symmetry in the position of these coordinates. `cell_shape_x` and `cell_shape_y` should define the position of all the vacancies given the coordinates for a single vacancy.
- `Bulk Structure`. Computing the grand potential (see paper eq. 6) relies on the bulk structure.
- `Mono Structure`. The monolayer defined without any adsorbate.
- `Calculation Parameters`. Parameters to define which charge values should be tested. Also sets the axis that the monolayer is defined (since this is a 2d object, and the axis will affect some of the automation). Sets the cell shape (x, y) which defines how many adsorbates can be added to the surface. `reflect_vacancies` is a boolean that determines whether we want to place adsorbates on both sides of the monolayer (which should almost always be the case. Vacancies should not define the positions on the other side, and they are used to determine the possible configurations. These adsorbates are then reflected since each configuration ought to be symmetric through the axis that the monolayer is defined).

`gen_supercell` is the CalcFunction that generates all the structures to be used in the workflow. We generate all possible structures for a given size ( $x, y$ ), and given number of adsorbates to fill ( $n$ ). Any structure that is equivalent to another via rotation (multiples of  $\pi/2$ ), reflection (through principal axes) and translation will be removed since the energies for these duplicate configurations will be the same. This is achieved by defining a rudimentary intranslatable coordinate system and then looping through all possible rotation/reflection

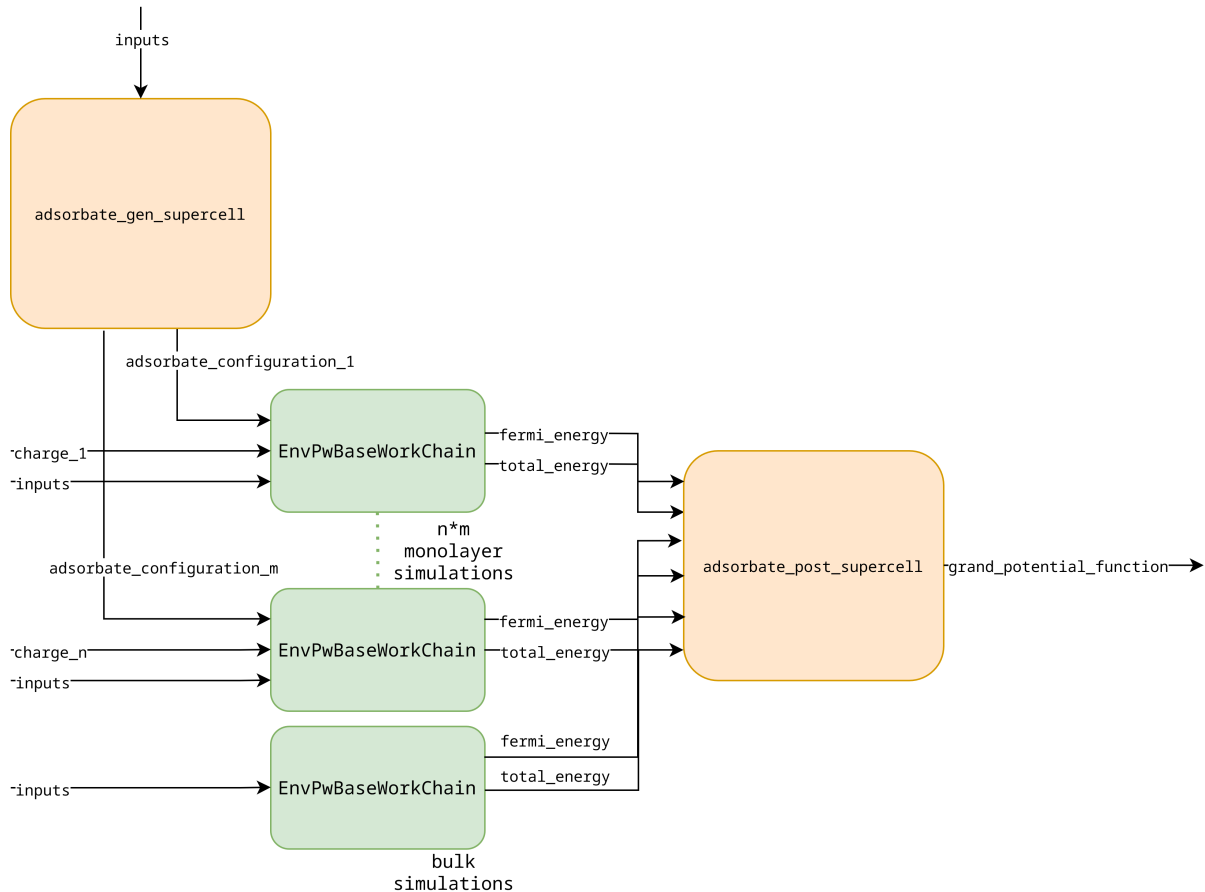


FIGURE 3.3. Grand canonical simulation workflow consisting of pre- and post-processing stages, and a series of monolayer and bulk simulations.

states.

For  $x$  charge increments and  $y$  possible adsorbate configurations, all the necessary simulations are performed. In total this equates to  $xy + 2$  simulations.

From each simulation, the total energy is taken, as well as the Fermi energy (and the correction to the Fermi energy from Environ). The goal is to compute the grand potential (see paper eq. 6). We populate three two-dimensional arrays, the free energy, the Fermi energy, and the free energy difference between the monolayer and the bulk. Each two-dimensional array contains values for a range of charge/adsorbate configurations.

`adsorbate_post_supercell` performs the postprocessing, the details of which are discussed in the paper.

### 3.5.3. Machine Learning Adsorption Energy

This workflow is designed to follow the graph based method of Ulissi and coworkers [101] to construct surface phase diagrams. The idea is to consider all possible surface configurations on a catalytic surface. Sites are chosen from chemical intuition and for each site, a number of adsorbates can be considered. The number of configurations for such a setup scales very quickly as one increases the number of possible sites and adsorbate species.

`adsorbate_gen_multitype` serves to automate the process of finding all the possible adsorbate permutations. The paper discusses how one can exploit the power of Gaussian Processes, a machine learning tool picked up by researchers in the field relatively recently, to minimize the number of simulations required to predict all possible configurations. The idea is that due to similarities between certain configurations, one can interpolate between two different structures in order to predict a configuration without running a simulation. Similarities are measured by considering how many adsorbates need to be added or removed from the system in order for two structures to be equivalent. By considering the configurations as nodes, and by defining adjacency as two nodes that are different by the addition or removal of a single adsorbate, one can build up a connected graph of all possible configurations.

It is subsequently reasonable to choose the maximally connected nodes as the subset to feed into the Gaussian Process model as training data.

## 3.6. Visualization

Visualization in science is a vitally important tool for translating data into a form that is intuitive for both scientists and the general public. In materials discovery, we typically wish to display how a theoretical material or compound might look atomically. For continuum solvation models, many quantities are described as fields. There are commonly two ways to visualize these, either by representing a scalar field as a cloud. Since the entire data spans some three dimensional space, a single screenshot cannot realistically convey all the information in that space, but if we can transform the data into some kind of range of values with reasonable spacing such that areas of interest have larger values than their surroundings, we can lower the opacity of smaller values so that they are more transparent, and increase the

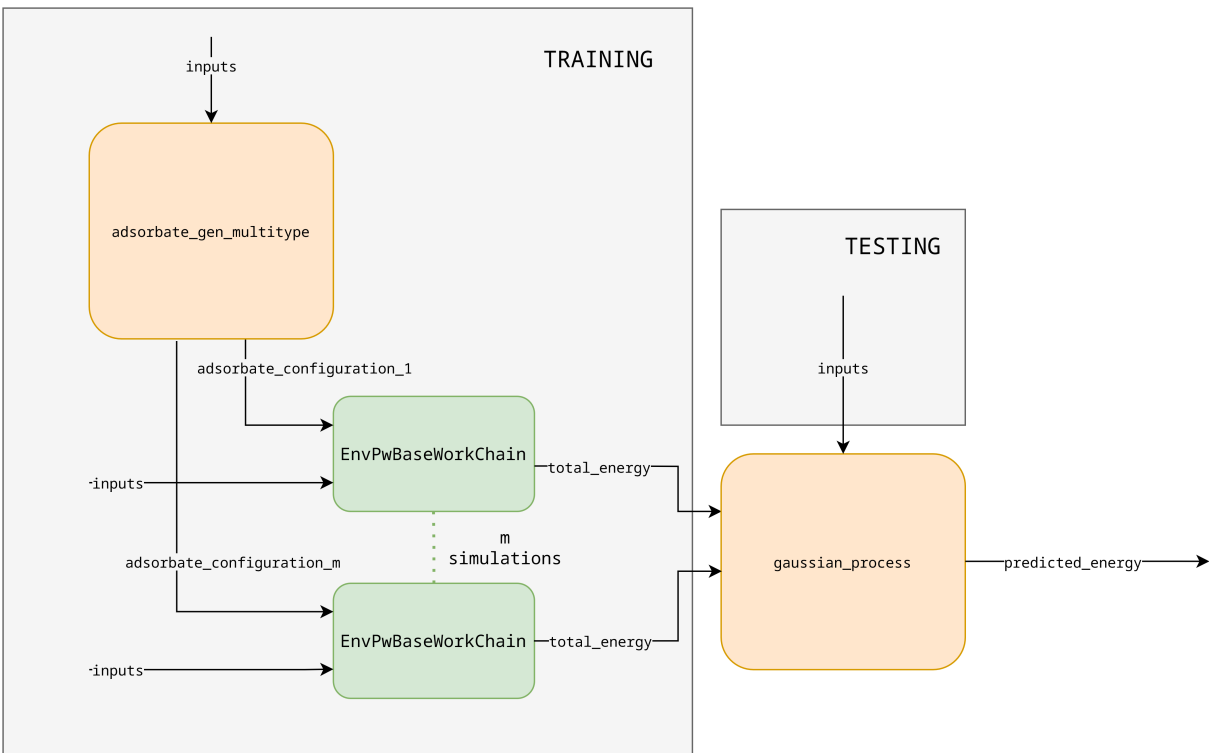


FIGURE 3.4. Graph based workflow that performs runs a minimal series of simulations to generate training data for adsorbate configurations.

opacity of large values such that they appear more opaque. The problem can be simplified by considering a single value within the range of values spanned by a scalar field. One can imagine generating an isosurface that represents all points that contain that value. This process typically assumes that the field is well behaved and there are no singularities. Since we have a discrete representation of the points in the field, interpolation is used to find the placement of the points that make up a 2d surface mesh.

An alternative representation for environmental boundaries (typically represented by isosurfaces) uses metaballs, which assumes the field object is atom-centered. Metaballs

Currently a handful of visualization tools exist, each of which sits on a scale from interactive simulation aids, to high quality image or animation rendering. The majority of my work sits on the latter of the scale, with the goal of providing researchers with a convenient tool for rendering images for journals and presentations. The user simply needs

an XML file defined by a simple schema that contains information on atomic positions, species, and bonds. The file can also contain render specific information, for example, any extra drawables, as well as light and camera positions and parameters. The tool itself is a plugin for the popular rendering tool Blender. Blender as of its 2.80 update in 2019 has gone from an open source alternative for enthusiasts and small teams on a budget to an impressive suite that rivals other leading software in the industry. I chose Blender since it was open source, and rapidly growing in popularity, as well as having a well-documented API for Python development. Blender as of 2019 already had an atomic add-on as part of its official collection of add-ons, however this tool only rendered atoms and bonds, and read in XYZ or PDB files directly.

Our requirements suggested a more ambitious project. A tool that worked with Quantum ESPRESSO and Environ output, and allowed not only for atom and bond rendering, but also the rendering of fields, as well as some more customization in the visualization options from the way that the atoms and bonds looked, to the materials used to represent them.

## CHAPTER 4

### CONTINUUM MODELS OF THE ELECTROCHEMICAL DIFFUSE LAYER IN ELECTRONIC-STRUCTURE CALCULATIONS

1

#### 4.1. Introduction

The electrical double layer (DL) is of primary importance in the field of energy conversion, as it plays a crucial role in devices such as supercapacitors and fuel cells[89, 13]. The DL structure is essentially characterized by two layers of opposite charge that appear at the interface between an electrified surface and an electrolyte solution. This structure arises from the charge accumulation at the boundary of the solvated surface, which attracts counterions from the bulk solution. The balance between the electrostatic attraction towards the charged surface, the entropic electrolyte contributions, and the steric repulsion between the ions gives rise to an equilibrium charge distribution in the solution that is generally known as the diffuse layer.

Unfortunately, various limitations hamper atomistic simulations of the diffuse layer[106]. First, long simulation times are required in order to achieve statistically significant samplings of the solvent and electrolyte configurations, with the corresponding time-scales being often beyond the reach of standard first-principles molecular dynamics techniques. In addition, large simulation cells are necessary in order to capture the long-range screening of typical values of the surface charge densities.

Continuum models represent an attractive alternative to fully-atomistic models of electrolyte solutions. A continuum description of the solvent and of the ions allows, in fact, to bypass the computationally-intensive configurational sampling of the solution's degrees of freedom. In particular, our focus here is on hybrid methods, where a first-principles modeling of an electrified surface is coupled to a continuum description of the solution (Figure 4.1).

---

<sup>1</sup>This chapter is reproduced from Francesco Nattino, Matthew Truscott, Nicola Marzari, and Oliviero Andreussi. Continuum models of the electrochemical diffuse layer in electronic-structure calculations. *J. Chem. Phys.* 150, 041722 (2019). DOI: <https://doi.org/10.1063/1.5054588>, with permission from AIP Publishing




FIGURE 4.1. Visualization of the atomistic details of a typical electrochemical setup (A): The metal slab (silver spheres) is in contact with a neutral aqueous solution (oxygen in red, hydrogen in white) containing electrolyte species (cyan and yellow spheres). Continuum models are obtained by integrating out the atomistic degrees of freedom of the mobile species (e.g. water and electrolytes) and replacing them with homogeneous continuum bodies, whose boundaries reflect the physical separation between the QM system and the environment. In B the solvent boundary (green transparent surface) is reported together with the continuous charge density of the electrolyte (blue transparent field), as computed for a charged substrate. Different onset for the solvent and electrolyte continuum models can be imposed in the definition of the model or can be obtained by including additional repulsive interactions between the continuum electrolyte and the QM substrate.

These models are particularly appealing for the accuracy and predictive power that they can potentially have, as the processes occurring at or within the metal surface are described at a quantum-mechanical level, while the electrostatic screening of the diffuse layer is accounted for at a mean-field level.



Starting from highly simplified models of the double layer, which consist of a countercharge plane at a fixed distance from a charged metal surface[34, 59, 15], more complex diffuse layer models have been subsequently proposed and integrated into periodic density-functional theory (DFT) codes[71, 51, 25, 24, 26, 57, 42, 95, 94, 64]. However, despite the large variety of electrolyte models proposed, there is no consensus on the model features required to achieve a physically-sound description of the diffuse layer. On one hand, full-continuum models that are based on the solution of some form of the size-modified Poisson-Boltzmann (PB) equation have been shown to qualitatively or semi-quantitatively describe experimental data[11, 69, 10]. On the other hand, recent work from Sundararaman et al.[94] has suggested that non-linear effects in the dielectric continuum also play an important role, and should thus be accounted for in order to reproduce measured trends.

In this work, we tackle these issues by systematically analyzing the performance of a hierarchy of continuum diffuse layer models of increasing complexity. We show how the various models can be derived from similar expressions of a free-energy functional and how they can be implemented in the framework of DFT. We choose the differential capacitance (DC) of a model metal surface as a prototypical observable to compare and contrast the various electrolyte models. In particular, we consider a Ag(100) surface in an aqueous solution as study system, motivated by the availability of accurate experimental data [102] that have been widely used in the literature to validate diffuse layer models [11, 95, 94].

Results show that a size-modified Poisson-Boltzmann model is able to qualitatively capture the main features of experimental DC curves, including the minimum capacitance value at the potential of zero charge, and the two local maxima at higher and lower potentials. The choice of solvation cavity employed to separate the quantum-mechanical region from the continuum solvent region is also found to play an important role on the absolute value of the computed DC.

The article is structured as follows. Section 4.2 reviews the theoretical background of the diffuse layer models considered and it presents the details of their computational implementations. Results on the computed DC values for the Ag(100) surface are then presented

in Section 4.3. In particular, the various electrolyte models are illustratively compared under vacuum conditions in Section 4.3.1, while their performance is better validated in Section 4.3.2 by comparing results to experimental data. Finally, the conclusions are presented in Section 4.4.

## 4.2. Methods

### 4.2.1. The Electrolyte Cavity

In the framework of continuum solvation models, the solvent's degrees of freedoms are smeared out in a continuum description and accounted for by means of a dielectric continuum. An important element in this class of models is represented by the so-called solvation cavity, which defines the boundary between the explicitly described solute region and the solvent region, where the dielectric continuum is located. This partitioning of the simulation cell can be formally defined through an interface function,  $s(\mathbf{r})$ . We define here  $s(\mathbf{r}) \equiv 1$  inside the solute region, and  $s(\mathbf{r}) \equiv 0$  in the region of space characterized by the solvent dielectric constant  $\varepsilon_0$ . In the field of material science and condensed matter physics, continuum models typically involve interface functions that smoothly switch between the two regions, as they turn out to provide a considerable improvement to numerical stability[28, 86, 3]. Closely related to  $s(\mathbf{r})$ , the dielectric function  $\varepsilon(\mathbf{r})$  sets the local value of the dielectric constant:

$$(4.1) \quad \varepsilon(\mathbf{r}) = 1 + (\varepsilon_0 - 1)(1 - s(\mathbf{r})).$$

In a similar fashion, the interface function  $s(\mathbf{r})$  can be exploited to define the region of space that is accessible to the ionic species in the solution. In particular, the complementary interface function  $\gamma(\mathbf{r})$  defines the portion of the cell where the electrolyte solution is located:

$$(4.2) \quad \gamma(\mathbf{r}) = 1 - s(\mathbf{r}).$$

It is important to stress here that in the above equations we have expressed both the solvent and the electrolyte domains in terms of the same interface function. In principle,

since the two domains are associated with the regions of space that are accessible to solvent and electrolyte particles, respectively, different interface functions should be required. For example, electrolyte ions that have strong solvation shells may be hindered direct access to the electrochemical interface, their closest distance from the substrate being increased by the thickness of the coordinating solvent molecules in what is known as the Stern layer[93]. In order to simplify the parameterization and tuning of the different interfaces, the electrolyte boundary is often expressed as a scaled version of the solvent one[24, 44, 81, 94]. Alternatively, a single interface is used and additional repulsive potentials are added to the free energy functional of the electrolyte system to stabilize solutions which are displaced from the solvent interface[51]. The latter approach has the additional flexibility of allowing both repulsive and attractive interactions between the components of the diffuse layer and the substrate. For this reason, we decided to focus the following discussion on models with a single common interface function.

The interface function  $s(\mathbf{r})$  is typically constructed as a function of the solute’s degrees of freedom. In order to explore how the choice of the solvation cavity affects the resulting diffuse layer model, we have considered the following three interface functions.

A first interface function is based on the local value of the solute’s electron density. An optimally-smooth switching function has been proposed by Andreussi et al.[3], who have revised in the so-called self-consistent continuum solvation (SCCS) model the one originally proposed by Fattbert and Gigy[28] and Scherlis et al. [86].

As a second interface function, we consider a rigid function of the solute’s atomic positions, as defined through the product of atom-centered interlocking spheres with a smooth erf-like profile. This cavity function, as proposed by Fiscaro et al. in the recent soft-sphere continuum solvation (SSCS) model[31], accounts for the diversity of the chemical species involved through tabulated van-der-Waals radii.

The last interface function reflects the two-dimensional character of the slab system considered. In particular, a planar boundary between solvent and solute is defined as a mere function of the vertical distance from the slab center,  $d$ . Two parameters are employed to

define the cavity. The first one,  $d$ , defines the absolute position of the interface. The second one,  $\Delta$ , regulates the smoothness of the boundary by defining the spread of the error-function profile which we have included along the surface normal. This simplified form of interface has not been explicitly addressed before in the literature. It has been introduced here to facilitate the comparison with analytical one-dimensional models. For particularly regular substrates it may provide an easier and more robust approach with respect to atomic-centered or electron-based interfaces.

#### 4.2.2. The Electrolyte Models

In this section we describe the continuum electrolyte models considered and illustrate how they can be derived from specific free-energy functionals where we include all electrostatic and mean-field contributions (the usual non-interacting electron kinetic energy and the exchange-correlation energy terms are also included, but left out in the expressions for improved readability). In order to perform self-consistent DFT calculations for a system embedded in an electrolyte solution, energy contributions that explicitly depend on the solute electron density require the inclusion of corresponding terms in the Kohn-Sham (KS) potential. Furthermore, terms that explicitly depend on the solute’s atomic positions give rise to analogous contributions to the atomic forces. All these contributions are reported in the Supplemental Material.

In this work, we have neglected the solvent-related non-electrostatic contributions to the free-energy[86, 3] based on the quantum volume and quantum surface[22]. Such contributions, however, can be straightforwardly included by adding corresponding terms to the free-energy, to the KS potential and to the forces[86, 3].

##### 4.2.2.1. Planar Countercharge Model

To first approximation, the electrolyte screening of the surface charge can be accounted for by introducing a countercharge plane at a given distance from the surface[34]. The presence of this external charge modifies the electrostatic energy of the system. When accounting for this electrostatic term, the free energy of the system embedded in the elec-

trolyte solution can be computed as:

$$(4.3) \quad F^{PC}[\rho(\mathbf{r}), \phi(\mathbf{r})] = \int \left[ -\frac{\varepsilon(\mathbf{r})}{8\pi} |\nabla\phi(\mathbf{r})|^2 + \rho(\mathbf{r})\phi(\mathbf{r}) + \rho^{ions}(\mathbf{r})\phi(\mathbf{r}) \right] d\mathbf{r}$$

Here  $\rho(\mathbf{r})$  is the total (i.e. electronic + nuclear) charge density of the solute,  $\phi(\mathbf{r})$  is the electrostatic potential and  $\rho^{ions}(\mathbf{r})$  is the external charge density that mimics the counterion accumulation.

This model can be seen as a computational implementation of the Helmholtz model for the double layer[46]: the countercharge plane completely screens the surface charge in a region of space that can be chosen to be infinitely narrow. Note that the Helmholtz screening does not depend on the ionic strength of the solution; it is thus not surprising that the bulk electrolyte concentration does not appear in Eq. 4.3.

#### 4.2.2.2. Poisson-Boltzmann Model

A more physical description of the diffuse layer can be derived from a free-energy expression that accounts for the chemical potential and the entropy of the ions in the solution[16, 24]. These terms allow one to introduce an explicit dependence on the local concentrations of the electrolyte species ( $\{c_i(\mathbf{r})\}$ ). For an electrolyte solution with  $p$  ionic species with charges  $\{z_i\}$  and bulk concentrations  $\{c_i^0\}$ , such that the solution is overall neutral ( $\sum_{i=1}^p z_i c_i^0 = 0$ ), the free energy functional takes the following form[16, 24]:

$$(4.4) \quad F[\rho(\mathbf{r}), \phi(\mathbf{r}), \{c_i(\mathbf{r})\}] = \int \left[ -\frac{\varepsilon(\mathbf{r})}{8\pi} |\nabla\phi(\mathbf{r})|^2 + \rho(\mathbf{r})\phi(\mathbf{r}) + \rho^{ions}(\mathbf{r})\phi(\mathbf{r}) + \right. \\ \left. - \sum_{i=1}^p \mu_i (c_i(\mathbf{r}) - c_i^0) - T (s[\{c_i(\mathbf{r})\}] - s[\{c_i^0\}]) \right] d\mathbf{r}.$$

Here  $\mu_i$  is the chemical potential of the  $i$ -th electrolyte species,  $T$  is the temperature, and  $s[\{c_i\}]$  is the electrolyte entropy density per unit volume. The electrolyte charge density can be expressed in terms of the local electrolyte concentrations as  $\rho^{ions}(\mathbf{r}) = \sum_{i=1}^p c_i(\mathbf{r}) z_i$ .

Under the assumptions of a point-charge electrolyte and ideal mixing, the entropy density of the solution is:

$$(4.5) \quad s[\{c_i(\mathbf{r})\}] = -k_B \sum_{i=1}^p c_i(\mathbf{r}) \ln \frac{c_i(\mathbf{r})}{\gamma(\mathbf{r})},$$

where  $k_B$  is the Boltzmann constant. Note that the exclusion function  $\gamma(\mathbf{r})$ , which sets the boundary between the electrolyte solution and the solute region, enforces a zero entropic contribution from the volume assigned to the quantum-mechanical region.

In order to find an expression for the equilibrium electrolyte concentrations, the free-energy functional in equation 4.4 is minimized with respect to  $c_i(\mathbf{r})$ . This procedure first leads to the condition:

$$(4.6) \quad z_i \phi(\mathbf{r}) - \mu_i + k_B T \left( \ln \frac{c_i(\mathbf{r})}{\gamma(\mathbf{r})} + 1 \right) = 0,$$

which allows one to obtain an expression for the chemical potential from the bulk electrolyte region, where  $\phi(\mathbf{r}) = 0$  and  $\gamma(\mathbf{r}) = 1$ , obtaining:

$$(4.7) \quad \mu_i = k_B T (\ln c_i^0 + 1).$$

By substituting equation 4.7 back into equation 4.6 one then obtains the following expression for the equilibrium electrolyte concentration:

$$(4.8) \quad c_i(\mathbf{r}) = \gamma(\mathbf{r}) c_i^0 e^{-\frac{z_i \phi(\mathbf{r})}{k_B T}} \equiv c_i^{PB}(\phi(\mathbf{r}))$$

By using this equilibrium electrolyte concentration, the free-energy functional expression significantly simplifies to:

$$(4.9) \quad F^{PB}[\rho(\mathbf{r}), \phi(\mathbf{r})] = \int \left[ -\frac{\varepsilon(\mathbf{r})}{8\pi} |\nabla \phi(\mathbf{r})|^2 + \rho(\mathbf{r}) \phi(\mathbf{r}) + k_B T \sum_{i=1}^p c_i^0 \left( 1 - \gamma(\mathbf{r}) e^{-\frac{z_i \phi(\mathbf{r})}{k_B T}} \right) \right] d\mathbf{r}.$$

Minimization of the free-energy functional with respect to  $\phi(\mathbf{r})$  now leads to the well-known Poisson-Boltzmann equation (PBE), which allows one to relate the equilibrium charge densities in the system to the electrostatic potential of the system:

$$(4.10) \quad \nabla \cdot \varepsilon(\mathbf{r}) \nabla \phi(\mathbf{r}) + 4\pi \sum_{i=1}^p z_i c_i^{PB}(\phi(\mathbf{r})) = -4\pi \rho(\mathbf{r}).$$

For low electrostatic potentials, i.e. whenever  $z_i \phi(\mathbf{r}) \ll k_B T$ , one can approximate the exponential dependence on  $\phi(\mathbf{r})$  with a linear function:

$$(4.11) \quad e^{-\frac{z_i \phi(\mathbf{r})}{k_B T}} \approx 1 - \frac{z_i \phi(\mathbf{r})}{k_B T}.$$

The expression of the electrolyte concentrations thus reduces to:

$$(4.12) \quad c_i(\mathbf{r}) \approx \gamma(\mathbf{r}) c_i^0 \left( 1 - \frac{z_i \phi(\mathbf{r})}{k_B T} \right) \equiv c^{LPB}(\phi(\mathbf{r})),$$

which leads to the following linearized-version of the PBE (LPBE):

$$(4.13) \quad \nabla \cdot \varepsilon(\mathbf{r}) \nabla \phi(\mathbf{r}) - k^2 \gamma(\mathbf{r}) \phi(\mathbf{r}) = -4\pi \rho(\mathbf{r}).$$

The constant operator  $k^2 = 4\pi \frac{\sum_{i=1}^p z_i^2 c_i^0}{k_B T}$  is related to the Debye length  $\lambda_D$  of the electrolyte solution:

$$(4.14) \quad k^2 = \frac{\varepsilon_0}{\lambda_D^2}.$$

The LPBE can be equivalently derived[81] by Taylor-expanding the exponential term in Eq. 4.9 up to second order in  $\phi(\mathbf{r})$ , and by subsequently minimizing the resulting energy functional,

$$(4.15) \quad F^{LPB}[\rho(\mathbf{r}), \phi(\mathbf{r})] = \int \left[ -\frac{\varepsilon(\mathbf{r})}{8\pi} |\nabla \phi(\mathbf{r})|^2 + \rho(\mathbf{r}) \phi(\mathbf{r}) + \right.$$

$$\left. -\frac{\sum_{i=1}^p z_i^2 c_i^0}{2k_B T} \gamma(\mathbf{r}) \phi^2(\mathbf{r}) + k_B T \sum_{i=1}^p c_i^0 (1 - \gamma(\mathbf{r})) \right] d\mathbf{r},$$

with respect to the electrostatic potential.

The linear-regime of the PBE is expected to hold for a narrow potential range around the potential of zero charge (PZC). However, typical applications easily require the modeling of potential windows that extend for hundreds of mV, making it desirable to have efficient strategies to solve the full PB problem instead. If the interface between the solute and the electrolyte is suitable to a two-dimensional approximation, one can tackle the full PB problem by taking advantage of the reduced dimensionality of the interface. In particular, one can integrate out the dimensions in the surface plane and exploit the analytical solution of the PBE in one dimension[24, 25, 26]. In the following, we assume for convenience that the system is oriented with the  $x$  axis perpendicular to the slab plane and that the diffuse layer starts at a distance  $x_{Stern}$  from the center of the slab. Taking the planar average of the physical quantities involved in Eq. 4.10 and assuming that the system charge density and the dielectric interfaces are fully contained within the  $x_{Stern}$  distance, the resulting one-dimensional differential equation

$$(4.16) \quad \frac{d^2 \phi(x)}{dx^2} = -\frac{4\pi}{\varepsilon_0} \left( \rho(x) + \sum_{i=1}^p z_i c_i(\phi(x)) \right)$$

can be integrated analytically for  $|x| \geq x_{Stern}$ . For the most common case of a diffuse layer composed by ions of equal concentrations  $c^0$  and opposite signs, the electrostatic potential in the electrolyte region can be expressed as (see Supplemental material for the derivation):

$$(4.17) \quad \phi(x) = \frac{4k_B T}{|z|} \coth^{-1}(c_1 \exp(c_2 |x|)),$$

where  $c_2 = 32\pi k_B T c^0 / \varepsilon_0$  and  $c_1$  is obtained by imposing continuity of the normal component of the electric field at the electrolyte interface (i.e. for  $x = x_{Stern}$ ). Only the solution with the correct asymptotic behavior has been selected, ensuring that  $\phi(|x| \rightarrow \infty) =$



0. This model effectively corresponds to the Gouy-Chapman (GC) model for the diffuse layer[41, 20], and shares the assumption of a planar distribution of charge at a fixed distance from the slab surface with the Helmholtz model described in Section 4.2.2.1. However, it includes a more physically sound shape of the diffuse layer along the surface normal. In the linearized regime, the one-dimensional solution of the electrostatic problem would instead be given by (see Supplemental material for the derivation):

$$(4.18) \quad \phi(x) = c_0 \exp\left(\frac{k}{\varepsilon_0} |x|\right)$$

where  $c_0$  can be obtained by imposing continuity of the normal component of the electric field at the electrolyte interface.

For complex interfaces and general geometries, and for applications for which the linear-regime of the PBE is not expected to hold, one needs to numerically solve the full non-linear PBE (Eq. 4.10) to find the electrolyte concentration that minimizes the energy of the solvated system.

#### 4.2.2.3. Size-Modified Poisson-Boltzmann Model

The standard PB model assumes point-like ions, and consistently overestimates the electrolyte countercharge accumulation at electrode surfaces. An improved model for the diffuse layer accounts for the steric repulsion between the ions, which opposes the electrostatic attraction towards the electrode surface and therefore limits electrolyte crowding. This is the so-called size-modified PB (MPB) model, which can be derived from the free-energy functional as in Eq. 4.4 but exploiting an entropy density expression that accounts for the finite-size of the ionic particles.

Borukhov et al. [16, 17] derived such an entropy expression from a lattice-gas model. In particular, the volume of the continuum solution is divided into a three dimensional lattice, with each cell of the lattice being occupied by no more than one ion. Thus, the cell volume  $a^3$  or, equivalently, the maximum local ionic concentration  $c_0 = \frac{1}{a^3}$ , sets the distance of closest approach between ionic particles in the solution. In this framework, the solute region is not

part of the continuum solution and should therefore give zero contribution to the solution entropy density. Otani and Sugino, who focused on two-dimensional slab systems, naturally achieved this limit by setting the boundary for the continuum solution region at a fixed distance from the surface[71]. In their derivations, Jinnouchi and Anderson [51] have instead imposed such a limit through an effective repulsive interaction between solute and electrolyte, which prevents the electrolyte solution from entering the quantum-mechanical region. Ringe et al. [81, 82] have similarly accounted for such repulsive potential by recasting it in the form of an exclusion function. Here we follow a different approach[24, 25], and enforce the limit by imposing a space-dependence for the maximum ionic concentration, consequently exploiting the complementary interface function  $\gamma(\mathbf{r})$ :  $c_0 \equiv c_0(\mathbf{r}) = c_{max}\gamma(\mathbf{r})$ . The final expression for the electrolyte entropy density is therefore:

$$(4.19) \quad s[\{c_i(\mathbf{r})\}] = -k_B \sum_{i=1}^p c_i(\mathbf{r}) \ln \frac{c_i(\mathbf{r})}{c_{max}\gamma(\mathbf{r})} + \\ - k_B \left( c_{max}\gamma(\mathbf{r}) - \sum_{i=1}^p c_i(\mathbf{r}) \right) \ln \left( 1 - \sum_{i=1}^p \frac{c_i(\mathbf{r})}{c_{max}\gamma(\mathbf{r})} \right),$$

The first and second terms in Eq. 4.19 can be identified as the entropy contributions from the ions and the solvent, respectively. As in Eq. 4.5, the exclusion function  $\gamma(\mathbf{r})$  sets the boundary of the region that contributes to the entropy of the electrolyte solution.

By minimizing the free-energy functional in Eq. 4.4 with respect to the ion concentration we obtain the following expressions for the electrolyte chemical potentials (cf. Eq. 4.7):

$$(4.20) \quad \mu_i = k_B T \ln \left( \frac{c_i^0}{c_{max} - \sum_{i=1}^p c_i^0} \right),$$

and for the equilibrium ionic concentrations (cf. Eq. 4.8):

$$(4.21) \quad c_i(\mathbf{r}) = \frac{\gamma(\mathbf{r}) c_i^0 e^{-\frac{z_i \phi(\mathbf{r})}{k_B T}}}{1 - \sum_{i=1}^p \frac{c_i^0}{c_{max}} \left(1 - e^{-\frac{z_i \phi(\mathbf{r})}{k_B T}}\right)} \equiv c_i^{MPB}(\phi(\mathbf{r}))$$

The denominator in Eq. 4.21 renormalizes the concentration in the regions where the electrostatic interaction energy is comparable to or larger than  $k_B T$ , and sets  $c_{max}$  as the maximum electrolyte concentration. This parameter can be related to the effective ionic radius  $r_i$  through  $c_{max} = \frac{3P}{4\pi r_i^3}$ . In the following, we will assume random close packing for the electrolyte particles and correspondingly set the packing efficiency  $P = 0.64$ . Note that the point charge limit of Eq. 4.21, which corresponds to  $c_{max} \rightarrow \infty$ , consistently leads to the equilibrium concentration as derived in the standard PB model (Eq. 4.8).

By substituting Eqs. 4.19-4.21 into Eq. 4.4, one obtains the following expression for the MPB free-energy functional:

$$(4.22) \quad F^{MPB}[\rho(\mathbf{r}), \phi(\mathbf{r})] = \int \left[ -\frac{\epsilon(\mathbf{r})}{8\pi} |\nabla \phi(\mathbf{r})|^2 + \rho(\mathbf{r}) \phi(\mathbf{r}) + k_B T c_{max} \gamma(\mathbf{r}) \ln \left( c_{max} - \sum_{i=1}^p c_i^0 \right) + \right. \\ \left. - k_B T c_{max} \gamma(\mathbf{r}) \ln \left( c_{max} - \sum_{i=1}^p c_i^0 \left( 1 - e^{-\frac{z_i \phi(\mathbf{r})}{k_B T}} \right) \right) \right] d\mathbf{r} - k_B T c_{max} V \ln \left( \frac{c_{max} - \sum_{i=1}^p c_i^0}{c_{max}} \right),$$

where  $V$  is the simulation cell volume. Minimization with respect to  $\phi(\mathbf{r})$  finally leads to the size-modified Poisson-Boltzmann equation (MPBE), which is analogous to the standard PBE (Eq. 4.10) where, however,  $c_i^{PB}(\phi(\mathbf{r}))$  is replaced by  $c_i^{MPB}(\phi(\mathbf{r}))$ .

#### 4.2.2.4. Additional Interactions

The MPB model accounts for the steric repulsion between the ions in the solution, which limits electrolyte crowding. In addition, the solute and the ionic particles are expected to be surrounded by a solvation shell, where diffusing electrolyte particles are not expected to enter. The presence of this solvent-accessible but ion-free region, generally known as the Stern layer[93], can be simulated in a continuum framework via finite spacing between the onset of the dielectric function and the electrolyte charge density.

Such a spacing can be effectively introduced through an *ad-hoc* repulsive term between solute and electrolyte,  $\varphi(\mathbf{r})$ [51]. The repulsive interaction would therefore give rise to the following free-energy contribution:

$$(4.23) \quad E_{rep}[\{c_i(\mathbf{r})\}] = \sum_{i=1}^p \int c_i(\mathbf{r}) \varphi(\mathbf{r}) d\mathbf{r},$$

and subsequently appear in the expression for the equilibrium electrolyte concentration:

$$(4.24) \quad c_i(\mathbf{r}) = \frac{\gamma(\mathbf{r}) c_i^0 e^{-\frac{z_i \phi(\mathbf{r}) + \varphi(\mathbf{r})}{k_B T}}}{1 - \sum_{i=1}^p \frac{c_i^0}{c_{max}} \left( 1 - e^{-\frac{z_i \phi(\mathbf{r}) + \varphi(\mathbf{r})}{k_B T}} \right)}$$

As noted by Ringe et al.[81], a repulsive solute-electrolyte interaction can be recast in the form of an electrolyte-specific exclusion function  $\alpha(\mathbf{r}) \equiv e^{-\frac{\varphi(\mathbf{r})}{k_b T}}$ . The exclusion function alone prevents the electrolyte from approaching and entering the solute region in their Stern-corrected MPB model. The exclusion function  $\gamma(\mathbf{r})$  that appears in our model has a different physical origin, since it reflects the hard separation between quantum-mechanical and continuum regions. While the presence of the Stern layer could be effectively included in our model by using separate interface functions for dielectric and electrolyte, as for instance done by Dabo et al.[24, 25], we find the picture of a single interface setting the boundary between quantum solute and continuum solution more physically sound. We thus resort to repulsive interactions to introduce the finite spacing between the onsets of the dielectric and the electrolyte fluids. Note that, similarly to Ringe's model, our approach also predicts a zero entropic contribution from the Stern-layer volume (cf. Eq. 4.19), which is consistent with the expected absence of diffusing solvent and electrolyte particles in this region.

For a two-dimensional system like a metal slab, we find appropriate to use one-dimensional exponential functions to define the repulsion potential:

$$(4.25) \quad \varphi(x) = e^{-\frac{|x-x_0|-d}{w}},$$

where  $x_0$  corresponds the  $x$  coordinate of the slab center and the parameters  $d$  and  $w$  set the position and decay rate of the potential, respectively.

Baskin and Prendergast [10] have proposed a similar formalism to account for the specific adsorption of electrolyte species. In particular, they have used a Morse-like potential in a fully-continuum model to mimic anion adsorption on the electrode surface:

$$(4.26) \quad \varphi(x) = E_{ads} \left( \left( 1 - e^{-\frac{|x-x_0|-d}{w}} \right)^2 - 1 \right),$$

where  $E_{ads}$  is the anion adsorption energy and  $d$  now defines the distance between the surface plane and the adsorbed anion species. We have tested the introduction of such an interaction term in our mixed first-principles-continuum model. This description is computationally attractive as it bypasses the need for surface configuration and adsorbate coverage samplings. However, it is clear that such a model cannot be expected to capture electronic-structure changes of the metal surface beyond mean-field electrostatic effects.

### 4.2.3. Computational Implementations

#### 4.2.3.1. Planar Countercharge Model

For all the models presented here, calculations are performed in a symmetric setup: the electrode surface is modeled by means of a two-dimensional slab exposing two identical faces to the continuum solution. The computational setup thus involves two metal-solvent interfaces. Two charge distributions are added in front of the outermost atomic layers to compensate for the net charge of the surface,  $q$ .

For numerical reasons, the sharp countercharge plane that characterizes the Helmholtz model for the diffuse layer is broadened to have a Gaussian-shaped profile along the surface normal direction  $x$ :

$$(4.27) \quad \rho_{ions}^{Helmholtz}(x) = \frac{q}{2A\sqrt{\pi}\Delta} e^{-\frac{(|x-x_0|-d)^2}{\Delta^2}},$$

where  $A$  is the surface area and the factor 2 at the denominator of the prefactor arises from the symmetric setup employed. The distance  $d$  from the slab center  $x_0$  and the spread parameter  $\Delta$  constitute the only two parameters in the model.

The countercharge distributions are straightforwardly added to the total charge of the system. The corresponding electrostatic potential does not require self-consistency, allowing for fast and stable simulations.

#### 4.2.3.2. Analytic Planar-Averaged Poisson-Boltzmann Model

In many ways, modeling the diffuse layer via the analytical one-dimensional solution to the Poisson-Boltzmann problem can be seen as a straightforward modification of the Helmholtz approach, in which the shape of the planar countercharges is no longer given by a Gaussian envelope of arbitrary spread, but rather obtained from a physically sound model. However, simply inserting the diffuse layer as a charge distribution in the simulation cell would incur significant numerical problems: the analytical results for the electrolyte concentrations can present very sharp features close to the interface, which cannot be described accurately with the standard numerical resolution of the electronic-structure calculation. Additionally, diffuse layers may have very long decaying length-scales, extending for tens of nanometers from the electrochemical interface, thus requiring large simulation cells. For these reasons, when possible a description in terms of the effects of the diffuse layer on the quantum-mechanical system, i.e. its electrostatic potential, is preferred.

The assumption behind the model is the one of two sharp interfaces at a fixed distance  $x_{Stern}$  from the slab center  $x_0$ , above and below the slab: the quantum-mechanical system is fully contained within the two interfaces, while the diffuse layer is fully in the outer regions and is uniform along the  $yz$  planes perpendicular to the slab normal. With this setup, the net effect of the diffuse layer on the system would be a uniform shift,  $\Delta\phi^{DL}$ , of the electrostatic potential in the quantum-mechanical region of space, provided that the latter is computed with open-boundary conditions (OBC) along the  $x$  axis. Thus, a possible definition of the full potential in the simulation cell is represented by

$$(4.28) \quad \phi(\mathbf{r}) = \begin{cases} \phi^{OBC}(\mathbf{r}) + \Delta\phi_0^{DL} & |x - x_0| < x_{Stern} \\ \phi^{DL}(x) & |x - x_0| \geq x_{Stern} \end{cases},$$

where  $\phi^{DL}$  is given by Eqs. (4.17) or (4.18), and the shift is computed as

$$(4.29) \quad \Delta\phi_0^{DL} = \phi^{DL}(x_{Stern}) - \langle \phi^{OBC} \rangle_{yz}(x_{Stern}).$$

Since the quantum-mechanical system is not bound to be perfectly homogeneous in the  $yz$  plane, its planar average,  $\langle \phi^{OBC} \rangle_{yz}(x) = A^{-1} \int \int \phi^{OBC}(\mathbf{r}) dydz$ , is used in the above equation, where  $A$  is the surface area in the simulation cell. As a consequence, discontinuities can be present in the potential of Eq. (4.28) when passing through the  $x_{Stern}$  interfaces. Even though these discontinuities happen in a region of space which is not occupied by the quantum-mechanical system, they may be a source of numerical instabilities. To overcome this limitation, a slightly different path can be followed, where the diffuse-layer contribution to the potential is expressed as a one-dimensional continuous and smooth correction defined in the whole simulation cell, namely

$$(4.30) \quad \phi(\mathbf{r}) = \phi^{OBC}(\mathbf{r}) + \Delta\phi^{DL}(x)$$

where

$$(4.31) \quad \Delta\phi^{DL}(x) = \begin{cases} \Delta\phi_0^{DL} & |x - x_0| < x_{Stern} \\ \phi^{DL}(x) - \langle \phi^{OBC} \rangle_{yz}(x) & |x - x_0| \geq x_{Stern} \end{cases}.$$

The planar average of the potential on the right-hand side of the above equation can then be approximated by the one-dimensional potential of a planar-averaged charge distribution, namely

$$(4.32) \quad \langle \phi^{OBC} \rangle_{yz} (|x - x_0| > x_{Stern}) \approx \langle \phi^{OBC} \rangle_{yz} (x_{Stern}) - \frac{2\pi q}{A\epsilon_0} |x - x_{Stern}|,$$

where  $q$  is the total charge of the quantum-mechanical system and we have used the classical electrostatics result for the potential of a planar charge distribution. With the above formulation, the correction and its first derivative are both continuous at the interface.

While the correction described here is similar in spirit to the ‘electrochemical boundary conditions’ from Refs. [25, 24, 26], our approach makes use of the electrostatic potential that analytically solves the PBE in order to determine the diffuse-layer contribution to  $\phi(\mathbf{r})$ , without the need of an iterative procedure. Another novel element of the procedure described is the implementation of the correction that corresponds to the linear-regime version of the PB problem, which allows for the validation of the corresponding numerical approach.

While the above correction is defined for open-boundary conditions, standard electronic-structure simulations usually exploit periodic-boundary conditions. In this case, an alternative expression of the electrostatic potential of the electrochemical interface can be obtained, which incorporates the handling of PBC artifacts and of the diffuse layer into a single continuous and smooth one-dimensional correction. In particular, an approximate OBC potential is obtained as

$$(4.33) \quad \phi^{OBC}(\mathbf{r}) \approx \phi^{PBC}(\mathbf{r}) + \Delta\phi^{2D}(x)$$

and assuming a planar-averaged charge distribution a parabolic correction can be expressed as

[5]

$$(4.34) \quad \Delta\phi^{2D}(x) = \frac{\alpha_{1D}}{L_x} q - \frac{2\pi q}{V} x^2 + \frac{4\pi}{V} d_x x - \frac{2\pi}{V} Q_{xx},$$

where  $\alpha_{1D} = \pi/3$  is the Madelung constant of a one-dimensional lattice,  $V$  is the cell volume,  $q$ ,  $d_x$ , and  $Q_{xx}$  are the monopole, dipole, and quadrupole moments along the  $x$  axis



of the charge distribution. The corrected electrostatic potential can thus be easily expressed in terms of the PBC one as

$$(4.35) \quad \phi(\mathbf{r}) \approx \phi^{PBC}(\mathbf{r}) + \Delta\phi^{2D}(x) + \Delta\phi^{DL}(x).$$

The above approach can be adopted for any situation where an analytical one-dimensional solution to the electrostatic problem is available. In the description of the diffuse layer, both the Gouy-Chapman model and its linearized version have analytical solutions that can be inserted into the  $\phi^{DL}(x)$  term. Although based on a planar-average approximation, this class of correction approaches has significant advantages, in terms of speed and stability, when compared to more advanced numerical solutions of the electrostatic equations.

#### 4.2.3.3. Linearized Poisson-Boltzmann Model

In order to tackle the linear-regime version of the PB problem, we solve the corresponding differential equation using a preconditioned gradient-based method as proposed by Fiscaro et al. [30], which we will only summarize here. Briefly, we apply a conjugate-gradient algorithm to solve the LPBE:

$$(4.36) \quad \underbrace{(\nabla \cdot \varepsilon(\mathbf{r}) \nabla - k^2 \gamma(\mathbf{r}))}_{\mathbf{A}} \phi(\mathbf{r}) = \underbrace{-4\pi\rho(\mathbf{r})}_{b(\mathbf{r})},$$

using the following preconditioning operator:

$$(4.37) \quad \mathbf{P} = \sqrt{\varepsilon(\mathbf{r})} \nabla^2 \sqrt{\varepsilon(\mathbf{r})}$$

Therefore, instead of minimizing the residual function  $r(\mathbf{r}) = b(\mathbf{r}) - \mathbf{A}\phi(\mathbf{r})$ , one finds the solution of the problem by minimizing the preconditioned residual  $v(\mathbf{r}) = \mathbf{P}^{-1}r(\mathbf{r}) = \mathbf{P}^{-1}(b(\mathbf{r}) - \mathbf{A}\phi(\mathbf{r}))$ . The algorithm has been proven to converge in a limited number of iterations for simple analytical cases[30]. In addition, the choice of the preconditioner minimizes

the computational effort required [30]. In fact, the action of the operator  $\mathbf{A}$  on the preconditioned residual  $v_n(\mathbf{r})$ , which needs to be computed at each of the  $n$ -th solver's iteration, can be efficiently estimated as:

$$\begin{aligned}
(4.38) \quad \mathbf{A}v_n(\mathbf{r}) &= (\nabla \cdot \varepsilon(\mathbf{r}) \nabla - k^2 \gamma(\mathbf{r})) v_n(\mathbf{r}) \\
&= (\varepsilon(\mathbf{r}) \nabla^2 + \nabla \varepsilon(\mathbf{r}) \cdot \nabla - k^2 \gamma(\mathbf{r})) v_n(\mathbf{r}) \\
&= - (q(\mathbf{r}) + k^2 \gamma(\mathbf{r})) v_n(\mathbf{r}) + r_n(\mathbf{r}),
\end{aligned}$$

where  $q(\mathbf{r}) = \sqrt{\varepsilon(\mathbf{r})} \nabla^2 \sqrt{\varepsilon(\mathbf{r})}$  and we have used  $r_n(\mathbf{r}) = \mathbf{P}v_n(\mathbf{r}) = \varepsilon(\mathbf{r}) \nabla^2 v_n(\mathbf{r}) + \nabla \varepsilon(\mathbf{r}) \cdot \nabla v_n(\mathbf{r}) + q(\mathbf{r}) v_n(\mathbf{r})$ . The term  $q(\mathbf{r})$  can be evaluated only once and stored in memory. Once the terms  $r_n(\mathbf{r})$  and  $v_n(\mathbf{r})$  are computed, the evaluation of  $\mathbf{A}v_n(\mathbf{r})$  at each of the following iterations requires only vector-vector multiplications. The bottleneck of the algorithm is represented by the calculation of  $v_n(\mathbf{r})$ , which is calculated as[30]:

$$(4.39) \quad v_n(\mathbf{r}) = \mathbf{P}^{-1}r_n(\mathbf{r}) = \frac{1}{\sqrt{\varepsilon(\mathbf{r})}} (\nabla^2)^{-1} \left( \frac{r_n(\mathbf{r})}{\sqrt{\varepsilon(\mathbf{r})}} \right).$$

The overall algorithm performance is thus highly dependent on the solution of the standard Poisson problem, which in our case is carried out in reciprocal space. The term  $r_n(\mathbf{r})$  is computed instead from the knowledge of the residual function at the previous step  $r_{n-1}(\mathbf{r})$  and other quantities derived from the preconditioned residual  $v_n(\mathbf{r})$ [30], as typically carried out in conjugate-gradient approaches.

While the algorithm from Fisticaro *et al.* was tested in Ref.[30] by solving the LPBE only for simple analytic potentials, we investigate here for the first time the performance of such preconditioned conjugate-gradient algorithm for realistic electrified interfaces through our novel implementation in the ENVIRON module[2] Quantum ESPRESSO[39, 37].

#### 4.2.3.4. Standard and Size-Modified Poisson-Boltzmann Model

The preconditioned conjugate gradient algorithm from Ref. [30] can only tackle linear problems, like the one represented by the linearized-PB equation. Fisticaro *et al.* have also

proposed an iterative algorithm devised to solve the full non-linear PB equation[30], which, however, turned out not to be sufficiently stable to deal with extended charged systems. For the numerical solution of the full (size-modified) PB equation, we thus resort to the more robust Newton-based algorithm as proposed by Ringe et al. [81], which we have also implemented in the development version of the ENVIRON module[2]. In particular, the free-energy functional minimization that leads to the PB equation is recast as a root-finding problem:

$$(4.40) \quad G[\phi(\mathbf{r})] = 0,$$

where  $G[\phi_n(\mathbf{r})]$  is the  $\phi(\mathbf{r})$  functional derivative of the free-energy  $F[\phi(\mathbf{r})]$ :

$$(4.41) \quad G[\phi_n(\mathbf{r})] = \nabla \cdot \varepsilon(\mathbf{r}) \nabla \phi_n(\mathbf{r}) + 4\pi \rho(\mathbf{r}) + 4\pi \sum_i^p z_i c_i(\phi_n(\mathbf{r})),$$

Following Newton's iterative algorithm, the estimate for  $\phi(\mathbf{r})$  at the the  $n$ -th step is obtained as:

$$(4.42) \quad \phi_{n+1}(\mathbf{r}) = \phi_n(\mathbf{r}) + \frac{G[\phi_n(\mathbf{r})]}{G'[\phi_n(\mathbf{r})]}.$$

Here  $G'[\phi_n(\mathbf{r})]$  is the Fréchet derivative of  $G[\phi_n(\mathbf{r})]$ :

$$(4.43) \quad G'[\phi_n(\mathbf{r})] = \nabla \cdot \varepsilon(\mathbf{r}) \nabla + 4\pi \sum_i^p z_i \frac{\partial c_i}{\partial \phi}(\phi_n(\mathbf{r})).$$

By rearranging the terms in Eq. 4.42, the following linear problem is recovered:

$$(4.44) \quad \left( \nabla \cdot \varepsilon(\mathbf{r}) \nabla + 4\pi \sum_i^p z_i \frac{\partial c_i}{\partial \phi}(\phi_n(\mathbf{r})) \right) \phi_{n+1}(\mathbf{r}) = \\ - 4\pi \left( \rho(\mathbf{r}) + \sum_i^p z_i c_i(\phi_n(\mathbf{r})) - \sum_i^p z_i \frac{\partial c_i}{\partial \phi}(\phi_n(\mathbf{r})) \phi_n(\mathbf{r}) \right).$$

Overall, the algorithm proceeds as follows: starting from an initial guess for  $\phi(\mathbf{r})$ , i.e.  $\phi_0(\mathbf{r})$ , the charge and screening terms in Eq. 4.44 are evaluated. The new guess for the electrostatic potential,  $\phi_1(\mathbf{r})$ , is then determined by solving the corresponding linear problem using the preconditioned conjugate-gradient procedure from Ref. [30], as described in Section 4.2.3.3. The charge and screening terms are updated, and the new linear problem solved to find  $\phi_2(\mathbf{r})$ . These steps are repeated until convergence is achieved.

#### 4.2.4. Computational Details

All the electrolyte models discussed have been implemented in the developer version of the ENVIRON module[2] for the Quantum ESPRESSO distribution[39, 37]. Differential capacitances have been calculated by numerically differentiating charge-potential curves. We have employed a canonical approach: we perform constant charge calculations and determine the applied potential  $U$  *a posteriori* from the difference between the asymptotic electrostatic potential and the Fermi energy of the system. Experimental data and simulations have been compared to each other using as potential reference the corresponding estimate of the potential of zero charge. In the simulations, this is the potential computed for a neutrally-charged surface,  $U_{PZC}$  (see Figure 4.2).

All calculations have been performed using the Perdew-Burke-Ernzerhof exchange-correlation functional[72] and pseudo-potentials from the GBRV set[35], which have been chosen according to guidelines from the Standard Solid-State Pseudopotential library[77] (SSSP efficiency 0.7). Cutoff energies for the plane wave and density expansions have been set to 35 Ry and 350 Ry, respectively. A  $\Gamma$ -centered 18x18x1 k-point grid (or equivalent) has been employed to sample the first Brillouin zone, using the cold smearing technique from Ref.[63] with smearing parameter  $\sigma = 0.01$  Ry.

The Ag(100) slab has been constructed using 8 atomic layers at the bulk equilibrium lattice constant (4.149 Å). The slab has been fully relaxed only in vacuum. While electrolyte- and solvent-related contributions to the atomic forces are accounted for and they in principle allow for fully self-consistent optimizations in the presence of the embedding continuum, test calculations for the present case show that further relaxations in implicit solvent negligibly

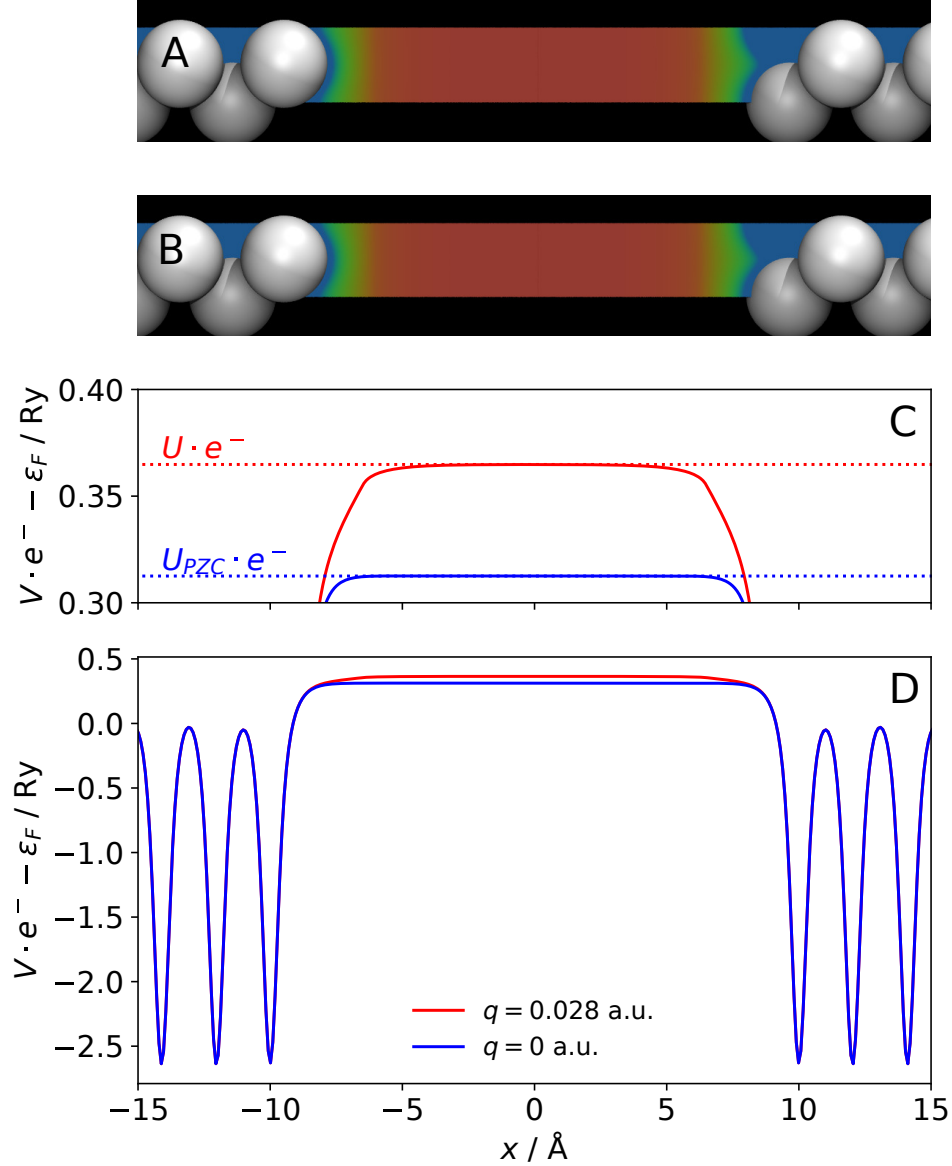


FIGURE 4.2. 2D-maps of the electrostatic potential computed with the planar-averaged analytical PB model (A) and the corresponding numerical version (B) for a Ag(100) slab in vacuum with a total charge  $q = 0.028$  a.u.. The planar-averaged electrostatic potential computed for the neutral and the charged surface (red and blue, respectively) is plotted as a function of the surface normal direction  $x$  in C and D. The same data is plotted in the C and D panels, using different scales for the potential axis. The planar interface with a symmetric monovalent electrolyte with bulk concentration  $c_0 = 0.1$  M has been employed.

affect the surface structure and the Fermi energy of the system.

Particular care needs to be taken with respect to the simulation cell size when using the SSCS (‘soft-sphere’) cavity. If the diameter of the atom-centered spheres exceeds one of the cell lattice vectors, a sharp transition arises in the region where the spheres overlap with their neighboring periodic replicas, which can trigger numerical instabilities. In order to obtain a smooth interface function, one has to set the simulation cell size such that all the sphere diameters fit the simulation box. For this reason, calculations using the SSCS have been performed in a (2x2) supercell. We also note that the Ag sphere radii as part of the original SSCS parameterization[31] are such that the resulting cavity includes non-physical dielectric pockets inside the metal slab. This issue has been fixed by introducing a non-local correction based on the convolution of the interface function with a solvent-size-related probe function. In this way, dielectric pockets that are smaller than the chosen solvent molecule (in this case water) can be identified and removed. Further details on the construction of such non-local interface are deferred to a forthcoming publication [1].

The models characterized by a self-consistent optimization of the ionic countercharge density require large separations between periodic replicas of the slab along the surface normal. This is to account for the long-range electrolyte screening, whose typical length is the Debye length  $\lambda_D$ . Due to the partial screening of the electrolyte charge by the dielectric, calculations that include implicit solvent require larger cell sizes along the surface normal. We have verified that Fermi energies are converged within few meVs for a 20 Å (60 Å) separation between periodic images for calculations in vacuum (implicit solvent). For calculations involving particularly low bulk ionic concentrations ( $\leq 0.04$  M) we have doubled these spacings. Note that the planar-averaged implementation of the PB model does not require these large spacings, as one resorts to the analytical solution of the one-dimensional problem to set the electrostatic potential at the cell boundaries. Such calculations have been thus performed with a spacing of 20 Å between periodic images. For both the numerical and analytic electrolyte models, we have made use of the parabolic corrective scheme from Ref. [5] in order to recover the potential of the isolated system from the electrostatic

potential computed with periodic-boundary conditions (see also Equations 4.33 and 4.34). This correction guarantees that the electrostatic potential approaches zero at large distances from the metal slab, provided that enough empty space is included in the unit cell for the numerical models. While charge neutrality is often enforced by means of a Lagrange multiplier  $\mu_{el}$  in the search for the electrostatic potential  $\phi(\mathbf{r})$  that minimizes the energy of the system [96, 42, 64], the asymptotically-zero reference potential can be used with  $\mu_{el} \equiv 0$ . This choice simultaneously provides the correct asymptotic limits for the electrolyte charge density and for the ionic concentration profiles[64].

### 4.3. Results and Discussion

#### 4.3.1. Vacuum

We start by considering the differential capacitance (DC) of Ag(100) in a solution with the vacuum dielectric constant ( $\epsilon_0 = 1$ ), which allows us to disentangle the electrolyte effects from the role played by the dielectric medium.

We first consider the planar-countercharge Helmholtz model (see Sections 4.2.2.1 and 4.2.3.1), which represents the lowest-rung diffuse layer model among the ones considered here. Figure 4.3 illustrates how the two parameters in the model, i.e. the surface-countercharge distance  $d$  and the spread of the charge distribution  $\Delta$ , affect the computed DC. Overall, charge-potential curves are found to be close to linear for all tested parameter values. The DC predicted by the Helmholtz model is thus almost potential independent, with a small DC decrease for increasing potentials. This trend is consistent with the larger electron-density spilling at more negative potentials, which effectively reduces the distance between the surface and the fixed countercharge distribution. This simple capacitor model also explains the effect of the  $d$  parameter on the computed DC, as we observe an increase (decrease) of the DC value for an inward (outward) shift of the neutralizing counterion density. The broadening of the electrolyte charge density, as regulated by the  $\Delta$  parameter, instead has a negligible effect on the DC. This is expected, as the spread of the distribution only affects the field in the narrow region where the countercharge is located. In contrast, the electrostatic




FIGURE 4.3. The charge per surface atom (in a.u.) is plotted as a function of the potential (in V). The Helmholtz model has been used for all data sets, varying the  $d$  and  $\Delta$  parameters that define the position and the width of the countercharge density, respectively. Note that the three lines corresponding to  $d = d_0/2 + 6$  Bohr are essentially superimposed ( $d_0$  is the slab thickness). The inset shows the differential capacitance as a function of the potential, as calculated from the analytical derivative of the spline-interpolated charge-potential curves (same line styles as in the main plot).

potential at large distances from the countercharge planes is essentially unaffected by the  $\Delta$  parameter, and so is the DC.

As already mentioned in Section 4.2.2.1, the Helmholtz model for the diffuse layer does not include any dependence on the bulk electrolyte concentration. The planar-averaged PB model overcomes this limitation while retaining the assumption of a planar countercharge-



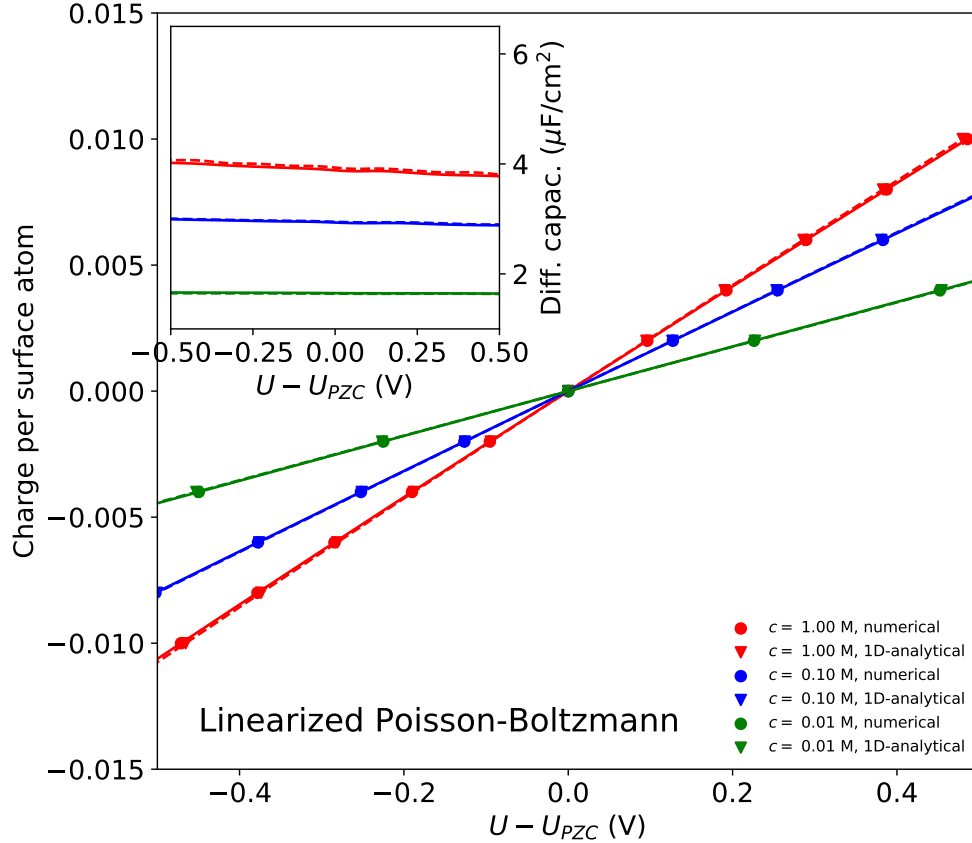


FIGURE 4.4. Same as Figure 4.3, but for the planar-averaged analytic linearized-PB model (triangles and dashed lines) and the corresponding numerical implementation (circles and solid lines). Red, blue and green symbols correspond to bulk electrolyte concentrations  $c_0 = 1$  M,  $c_0 = 0.1$  M and  $c_0 = 0.01$  M, respectively.  $x_{Stern}$  in the analytic model (Section 4.2.3.2) and  $d$  in the planar interface employed in the numerical model (Section 4.2.3.3) are set so that the interface lies in both cases 6.568 Bohr away from the surface.

density profile. Results obtained with the linearized-PB model (see Sections 4.2.2.2 and 4.2.3.2) are illustrated in Figure 4.4, which shows the computed charge-potential curves and corresponding capacitance values for three representative electrolyte concentrations. The linear-regime PB model predicts a weak potential dependence of the DC, as also observed for the Helmholtz model, but the computed capacitance now depends on the electrolyte concentration. In particular, lower DC values correspond to lower ionic concentrations.

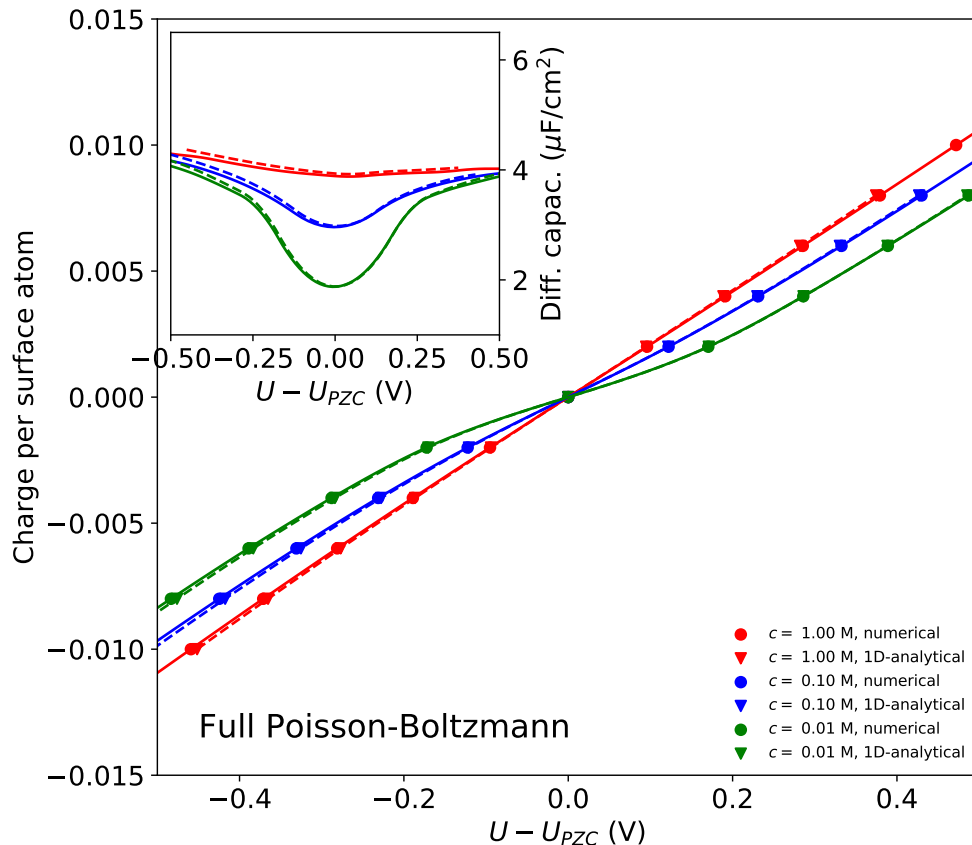


FIGURE 4.5. Same as Figure 4.4, but for the full PB model, in the planar-averaged analytic implementation (triangles and dashed lines) and the numerical implementation (circles and solid lines).

Figure 4.4 also includes results of calculations performed with the numerical linearized PB solver (see Sections 4.2.2.2 and 4.2.3.3), using a planar but smooth interface function with an error-function profile along the surface normal. We have used here a small spread parameter (0.01 Bohr) in order to better compare results to the planar-averaged LPB model, in which a sharp planar interface defines the boundary of the region where the one-dimensional LPBE is analytically solved. The DC computed through the numerical solution of the LPBE agrees well with what obtained through the corresponding planar-averaged analytic model. This is consistent with the interface being essentially two-dimensional, as expected for closely-packed metal surfaces like Ag(100).

The capacitance-potential trends obtained from the solution of the full-PBE (see Sec-

tions 4.2.2.2 and 4.2.3.4) are quite different, as illustrated in Figure 4.5. In contrast with the linear-regime model, the DC curves computed with the non-linear electrolyte model exhibit a concentration-dependent drop at the potential of zero charge (PZC), while similar capacitance values are observed at the largest potentials simulated for all electrolyte concentrations (see also Figs. 4-10 and 4-11 of Ref.[25]). As also observed for the linearized model, we find very good agreement between the capacitance curves computed with the planar-averaged approach, which exploits the analytical solution of the PBE along the surface normal, and the full numerical implementation. As illustrated in Figure 4.2, in fact, the electrostatic potential obtained with the numerical model is not significantly corrugated in the  $yz$ -plane at sufficiently large distance from the surface, and is thus very similar to the potential computed with the planar-averaged analytical approach.

The effect of the interface broadening on the DC is illustrated in Figure 4.6, where we compare DC-potential curves computed with the numerical solver of the full-PBE (Section 4.2.3.4) and a planar but smooth interface function. We have tested different values of the spread parameter  $\Delta$ , ranging from 0.01 to 0.5 Bohr. The DC is found to increase for increasing values of  $\Delta$ , which follows from the onset of the electrolyte-accessible region becoming closer to the surface. This effect is most pronounced at large (absolute) potentials and at high electrolyte concentrations. Under such conditions, in fact, the electrolyte charge density at the interface boundary is larger and sharper, and thus more sensitive to small changes in the onset region.

Figure 4.7 compares charge-potential curves and the corresponding DC values as computed with the numerical PB solver (Section 4.2.3.4) paired to the three different cavities presented in Section 4.2.1. Specifically, we have tested the use of the planar interface function also used in Figure 4.5 and 4.6 (see Figure 4.7A), and two additional cavities derived from the SSCS [31] (Figure 4.7B) and from the SCCS [3] (Figure 4.7C) models, respectively. To better compare results across the cavities employed, we choose the corresponding parameters so that the onsets of the three interface functions lie at approximately the same distance from the metal surface under neutral conditions, and a similar broadening characterizes the




FIGURE 4.6. Full Poisson-Boltzmann model: differential capacitance as a function of the potential for various values of the spread parameter  $\Delta$  in the planar interface function, using the numerical solver. The distance parameter  $d$  has been set so that the interface lies 6.568 Bohr away from the surface. The full numerical PB model has been used for all calculations, using  $c_0 = 0.01$  M.

three interfaces.

Very similar DC-potential curves are obtained using the planar and SSCS cavities. The former produces slightly higher capacitance values, consistently with the electrolyte charge density more closely approaching interstitial surface regions with the soft-sphere interface. Both interfaces predict DC-potential curves that are asymmetric around the PZC, with slightly larger DC values at negative potentials as compared to the corresponding positive values. This is again consistent with the effective separation between the surface and the ionic density onset becoming smaller at negative potentials due to the larger electron density spilling towards the rigid electrolyte interface. Interestingly enough, the trend observed with






FIGURE 4.7. Same as Figure 4.5, but for the full-numerical PB model (circle and solid lines) and linearized PB model (triangles and dashed lines). The interface functions employed are the following: the planar interface (top,  $d = d_0/2 + 6.568$  Bohr, where  $d_0$  is the slab thickness, and  $\Delta = 0.470$  Bohr); the soft-sphere interface (middle,  $r = 6.568$  Bohr,  $\Delta = 0.470$  Bohr); and the SCCS interface (bottom,  $\rho_{max} = 10^{-4}$  a.u. and  $\rho_{min} = 10^{-5}$  a.u.). The top subplot also includes results obtained with the Helmholtz model (black squares and dotted line) as a comparison. The same  $d$  and  $\Delta$  parameters used for the planar interface have been employed to set the Gaussian countercharge density.

the SCCS cavity is reversed. This density-dependent interface function, in fact, shifts the ionic density onset further away from the surface as the slab charge becomes more negative, effectively increasing the electrolyte-slab separation.

Figure 4.7 also includes results from the linearized PB model for the three cavities considered. This model correctly predicts the DC values at the PZC and the qualitative DC dependence on the bulk electrolyte concentration. Note that the capacitance computed with this model approaches the infinite-screening limit represented by the Helmholtz model capacitance for increasing ionic concentrations (Figure 4.7A). As also evident from comparing Figure 4.4 to Figure 4.5, the linearized version of the PB model dramatically fails in reproducing the potential trend computed with the corresponding non-linear model and it returns weak potential dependences with no minimum at the PZC. The monotonic trends observed for the linear-regime model are consistent with the patterns described for the corresponding non-linear model. For instance, the density-dependent SCCS cavity predicts monotonically increasing capacitance curves, as the surface-electrolyte gap increases with increasing potential. This is also consistent with the findings of Letchworth-Weaver and Arias[57], who have similarly observed monotonically increasing capacitance curves using a linearized-PB model for the diffuse layer and a density-dependent interface function.

Figure 4.8 compares DC curves computed using the standard (non-linear) PB model

to results from the size-modified PB model (see Sections 4.2.2.3 and 4.2.3.4). In particular, we test values for the  $c_{max}$  parameter that range from 300 M to 3 M, which correspond to effective ionic radii from 0.95 Å to 4.39 Å. Introducing a finite size for the ions affects the DC at large applied potentials: while the standard PB model predicts monotonically increasing capacitance values for increasing values of the applied potential, the size-modified model predicts the DC to first reach a maximum and then decrease as the potential deviates from the PZC. The DC maximum is reached at lower values of the potential for decreasing values of the  $c_{max}$  parameter. For  $c_{max} = 3$  M the DC-potential curve even changes concavity, and the PZC becomes the maximum. The observed DC decrease can be explained by the following arguments. In contrast with the standard PB-model, which allows for infinitely large electrolyte concentrations at the interface boundary, its size-modified variant imposes a maximum local ionic concentration,  $c_{max}$ . When this maximum concentration is locally reached, the steric repulsion between the ions pushes the ionic charge density towards the bulk solvent region, effectively increasing the separation between the surface and the electrolyte charge, giving rise to the observed DC decrease.

#### 4.3.2. Implicit Solvent

After having investigated the performance of the diffuse layer models in vacuum, we switch to simulations in implicit water and compare results to prototypical experimental data, presented in Figure 4.9. In particular, we have considered data reported by Valette[102] on the differential capacitance of Ag(100) in a KPF<sub>6</sub> electrolyte solution. Consistently with commonly-observed experimental trends, the DC exhibits a ‘camel-back’ shape, with the minimum indicating the PZC.

As also indicated by Valette, the common potential value at which the potential drop is observed across the various electrolyte concentrations and the rather symmetric shape of the DC curve around the PZC suggest a negligible anion adsorption in this electrolyte solution.

On the basis of the results presented in Section 4.3.1, only the MPB model is expected to qualitatively reproduce the experimental potential trend, with the capacitance drop at the




FIGURE 4.8. The differential capacitance is plotted as a function of the potential for various values of the  $c_{max}$  parameter in the MPB model. The planar interface function has been used for all calculations, with  $d = 6.568$  Bohr and  $\Delta = 0.470$  Bohr.

PZC and the DC saturation and decrease at large applied potentials. Instead, the Helmholtz and the linearized PB model fail to predict the capacitance minimum at the PZC, and the standard PB model predicts monotonically increasing capacitance curves. Concerning the cavity, the various interface functions have been found to give rise to overall similar capacitance values in vacuum, at least for parameterizations that lead to similar electrolyte charge distributions. The planar and SSCS cavities produce essentially identical results, and for this reason in the following we will consider only the latter, which better suits general interface geometries.

The DC computed with the MPB model and the SSCS interface including the dielectric continuum are plotted in Figure 4.9. We have used here the original SSCS cavity






FIGURE 4.9. The differential capacitance is plotted as a function of the potential. Experimental data from Ref.[102] are plotted as dashed lines. Results of MPB simulations using the SCCS cavity with the original parameterization from Ref. [3] are plotted as solid lines. The value of  $c_{max}$  is set to 20 M. Red is for  $c_0 = 0.1$  M and blue is for  $c_0 = 0.04$  M.

parameters[3], which have been fitted to a database of solvation energies of neutral molecules. We remind, however, that the non-electrostatic solvation terms have been neglected here. Calculations are performed for the experimental bulk electrolyte concentrations (0.1 M and 0.04M) and for the steric repulsion between ions through the  $c_{max}$  parameter, which we have initially set to 20 M. Assuming a random close-packing for the ions, this value of  $c_{max}$  corresponds to an effective ionic radius of approximately 2.33 Å. In comparison, experimental upper-bounds for the bare (non-solvated) ionic radii are 2.65 Å [60] and 2.42 Å[85], for  $K^+$  and  $PF_6^-$ , respectively.

As already observed in the vacuum environment, the MPB model predicts a DC drop

at the PZC, which is more pronounced for the lowest electrolyte concentration simulated ( $c_0 = 0.04$  M). This is in qualitative agreement with measurements. However, the overall absolute magnitude of the DC is severely underestimated, and the potential dependence computed is also much weaker than in experiments.

Figure 4.10 illustrates how the inclusion of additional solute-electrolyte interactions in the MPB model (see Sections 4.2.2.4) affects the computed DC (note the different scale of the  $y$  axis). Figure 4.10A shows the effect of including a solute-electrolyte repulsion potential. This potential introduces a gap between the onset of the dielectric continuum and the one of the electrolyte countercharge density. In particular, we test various values for the distance parameter  $d$  in the chosen functional form for the repulsive potential (see Eq. 4.25). By setting the spread parameter  $w$  to  $0.25 \text{ \AA}$  we ensure a fast decay of the exponential repulsion. The effect of introducing this Stern-layer gap is essentially a rigid shift of the DC curve. The observed capacitance decrease is consistent with the corresponding increase of the surface-electrolyte charge distance for increasing  $d$  values, and the small magnitude of the shift is related to the large dielectric constant that characterizes the region where the electrolyte charge is located.

Figure 4.10B illustrates how the DC is affected by anion adsorption as accounted through the continuum model of Baskin and Prendergast[10]. For the solute-anion interactions, we set the following values for the Morse-potential parameters:  $w = 0.5 \text{ \AA}$  and  $d = d_0/2 + 1.5 \text{ \AA}$ , where  $d_0$  is the slab thickness. The adsorption energy  $E_{ads}$  is varied in a range from  $3k_B T$  to  $9k_B T$  (i.e. from 80 meV to 230 meV at 300 K). Note that we simulate the asymmetric anion adsorption without accounting for solute-cation interactions. At the most negative potentials considered, where the electrostatic attraction of cations is much stronger than the imposed solute-anion interaction, the anion adsorption does not alter the computed DC. At the highest potentials simulated, the additional attractive interaction between surface and anions increases the electrolyte countercharge at the interface, thereby increasing the DC. At intermediate potentials, the anion adsorption shifts the DC minimum from the PZC towards negative potentials, where the electrostatic interaction compensates




FIGURE 4.10. The differential capacitance plotted as a function of the potential. All data refer to MPB simulations with the SCCS interface. The original parameterization from Ref. [3] has been employed, and  $c_0 = 0.1$  M and  $c_{max} = 20$  M. In the top panel, a repulsive potential between solute and electrolyte is introduced. Different colors correspond to different values of the  $d$  parameter, as indicated ( $w$  is set to  $0.25\text{\AA}$ ). In the bottom panel, a Morse-potential interaction between solute and anions is employed instead, with  $d = d_0/2 + 1.5\text{\AA}$  and  $w = 0.5\text{\AA}$ . Different colors correspond to different values of the  $E_{ads}$  parameter, as indicated.

the anion attractive potential.

The cavity parameterization is found to have a much larger influence on the absolute value of the computed DC. This is illustrated in Figure 4.11, where we plot the DC-potential curves calculated with the original interface parameterization that was optimized for neutral isolated systems and with the two parameterizations that have been later proposed[27] to best fit anion and cation solvation energies, respectively. The DC computed with the cation-specific parameterization does not significantly differ from the one obtained from the original parameterization. This is consistent with the very similar values for the cavity parameters  $\rho_{max}$  and  $\rho_{min}$  in the two fits. Significantly different cavity parameters were instead found to best fit the anion database, and we consistently observe a considerable difference in the resulting capacitance. In particular, the anion-specific parameterization is characterized by smaller cavities, and the reduced gap between the surface and the continuum fluids give rise to larger DC values. As illustrated in Figure 4.10A, the spacing between the surface and the electrolyte countercharge density has a rather contained effect on the absolute DC value. These findings suggest that the gap between the electrode surface and the dielectric polarization charge is thus the main responsible for the large DC Dependence on the cavity parameterization, as also suggested by Sundararaman et al. [94]. Consistently, Melander *et al.*[64], who have employed a dielectric cavity based on the van der Waals radii of the surface atoms, have found that increasing the atomic radii leads to significantly lower capacitance values.

It is evident that none of the three SCCS cavities described so far is able to describe well experimental data: the original SCCS parameterization and the cation refit underestimate the measured DC, while the anion parameterization overestimates it. This is not surprising, considering that all three parameter sets have been fitted to solvation energies of isolated systems. Figure 4.11A also includes a DC curve computed with cavity parameters that have been recently fitted[47] to reproduce the theoretical estimate of the absolute PZC of Pt(111). This last parameterization overall returns the best agreement with experimental data, even though it underestimates the measured DC in the potential region close to the

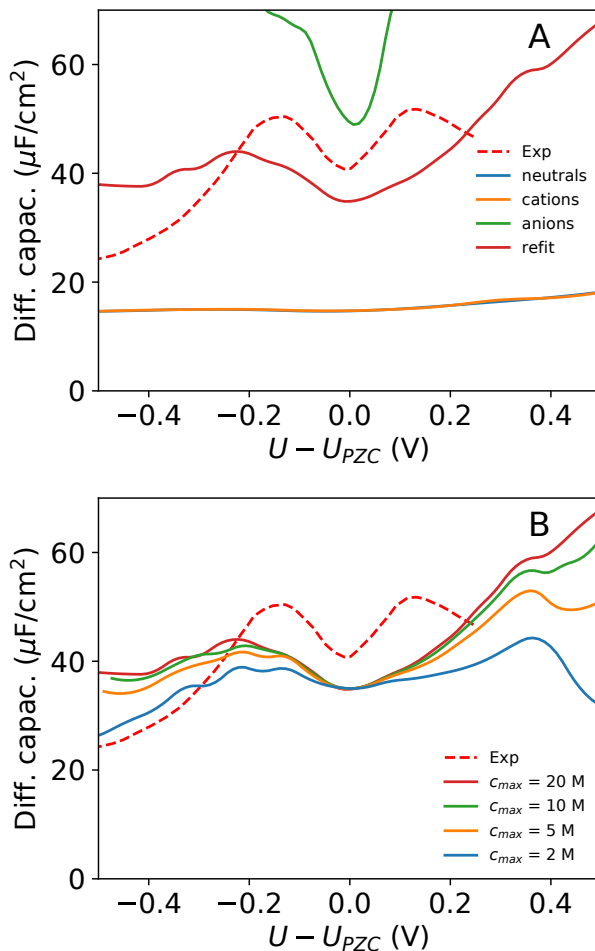


FIGURE 4.11. The differential capacitance is plotted as a function of the potential for  $c_0 = 0.1$  M. Experimental data[102] are shown as dashed lines. All theoretical data refer to MPB simulations with the SCCS interface function. In the top panel, results from the original SCCS parameterization [3] (‘neutral’), are compared to results from the cation- and anion-specific parameterizations[27] and to results from the parameter fit to the Pt(111) PZC[47]. The value of  $c_{max}$  is set to 20 M. In the bottom panel, the SCCS parameterization from Ref.[47] is employed, and the value of  $c_{max}$  varied from 20 M to 2 M, as indicated.

PZC and more severely overestimates it for larger applied potentials.

As clearly shown from the simulations in vacuum, the capacitance of the MPB model at large absolute potentials is strongly affected by the steric-repulsion between the ions (cf. Figure 4.8). Figure 4.11B shows the effect of decreasing the value of  $c_{max}$  from 20 M to 2 M, which is equivalent to increasing the ionic particle radii from from 2.33 Å to 5.02 Å. Decreasing  $c_{max}$  broadens the minimum in correspondence of the PZC and lower the capacitance at the highest and lowest potentials examined, improving agreement with experimental data. It is thus tempting to suggest that effective radii larger than the bare ones should be employed for the electrolyte particles, as also suggested in the literature on the basis of the strongly-bound solvent molecules surrounding ions in solution[11]. Note that experimental estimates for the radius of the solvated  $K^+$  ions span a range [11] from 3.8 Å [36] to 6.62 Å[70], which would largely justify the range of  $c_{max}$  investigated.

After having investigated the DC capacitance computed using density-based cavities, we now consider simulations performed with the SSCS interface function. Results are presented in Figure 4.12, where we have used the parameterization proposed by Fiscaro et al.[31]. Figure 4.12A reports the capacitance computed using the upper-bound  $c_{max}$  value of 20 M. Despite the cavity parameters were originally fitted to solvation energies of isolated systems as for the SCCS interface, the SSCS model leads to a very good description of the experimental DC around the PZC for the two electrolyte concentrations considered. As also observed for the SCCS interface function, the capacitance at large absolute potentials is instead overestimated when using  $c_{max} = 20$  M. Similar to the case of vacuum (Section 4.3.1), the (M)PB model returns asymmetric DC curves with the rigid cavity from the SSCS model. This finding can again be explained on the basis of the extent by which the electron density spilling from the metal surface approaches the continuum. The separation between the surface and the electrolyte onset, in fact, is effectively reduced for lower values of the potential, with a subsequent increase of the capacitance values.

Figure 4.12B shows the computed DC curve for the lower  $c_{max}$  value of 2 M ( $r_i = 5.02$  Å). The agreement with experiments is significantly improved, and both the position




FIGURE 4.12. Same as Figure 4.9, but the SSCS interface has been employed in the simulations. In top and bottom panel the value of  $c_{max}$  is set to 20 M and 2 M, respectively.

and the height of the ‘humps’ are comparable to measured data. Thus, also calculations performed with the SSCS cavity suggest that ionic radii larger than the bare ones should be employed to limit the steric crowding at electrode interfaces. Note that Sundararaman et al. recently achieved a similarly good description of the DC of Ag(100) with a soft-sphere-based continuum model[94]. The cavity size in their model was also based on the Ag ionic radius as tabulated in the unified force-field (UFF)[78], times a scaling constant. As also noted by Sundararaman et al. [94], the good description of the Ag(100) DC might be thus inferred to the Ag UFF ionic radius being fortuitously suitable to describe the cavity size for this system. Future investigations of the DC for other systems will shed light on this point.

Despite the overall good agreement with experimental data, the DC predicted using the SSCS cavity overestimates and underestimates the measured data at negative and positive potentials, respectively. The trend observed with the soft-sphere cavity in implicit solvent is consistent with the trends observed in vacuum with rigid cavities, which are found to predict an overall decreasing DC with increasing potential. In comparison, measurements exhibit slightly larger capacitance values at positive potentials, in better agreement with trends observed with the density-dependent SSCS cavity. Future work will clarify whether improved agreement with experimental data can be achieved by a specific refitting of the SSCS cavity.

Regardless on the cavity employed, our findings suggest that the MPB model for the electrolyte is able to capture the main features of the experimental DC for Ag(100) in an ideally non-adsorbing ionic solution. We note in passing that in addition of being more physically sound, the MPB model is also more numerically stable than the standard PB model, as the extremely large ionic charge densities that the latter predicts at the boundary between the electrified surface and the solvent region are difficult to handle with the numerical solvers without the inclusion of a Stern layer. While avoiding such instabilities, the linearized-PB model is inadequate for describing the capacitance of a charged metal surface. As expected from the results in vacuum (Section 4.3.1), Figure 4.13 illustrates how the linear-regime model predicts essentially potential-independent capacitance values, which are only accurate






FIGURE 4.13. Same as Figure 4.12, but the DC curves obtained with the MPB model and  $c_{max} = 20$  M are compared to analogous curves obtained with the linearized PB model.

close to the PZC.

Our results are in contrast with the finding from Ref. [94], where an additional non-linear dielectric model was suggested to be necessary in order to reproduce the trends observed in the measurements. Our findings also differ from the ones of Melander *et al.*[64], who have reported potential-independent capacitance trends for a metal surface (Au(210)) in an electrolyte solution using both a linearized- and a non-linear- MPB model. While discrepancies with Refs. [94, 64] require further investigation, our results are consistent with data from full-continuum models [11, 69, 10], where the solution of the non-linear MPBE is found to lead to the experimentally-observed ‘camel-back’-shape for the DC curve for metal surfaces in aqueous solutions.

#### 4.4. Summary and Conclusions

In summary, we have presented a hierarchy of electrolyte models that can be integrated in the framework of DFT to account for the presence of the diffuse layer in first-principles simulations of electrochemical interfaces. We have validated the accuracy of the models by comparing computed DC values to experimental data, focusing on the Ag(100) surface in an aqueous electrolyte as study system.

Results suggest that the size-modified PB model is necessary in order to reproduce the main characteristics of the experimental DC, i.e. the concentration-dependent drop at the PZC and the two local maxima at intermediate applied potentials. The lowest-rung planar Helmholtz model, which does not include any dependence on the bulk electrolyte concentration, predicts negligible DC dependences on the applied potential. Similarly, the standard PB model, both in the linear-regime and in its full non-linear implementations fails in describing experimental DC trends.

Further accounting for solvent effects through a continuum dielectric allows for a direct comparison of computed DC values to experimental data. We observe a large influence of the choice of the dielectric cavity on the absolute DC values, consistently with previous findings [95, 94]. For the SCCS interface function, the best agreement with experimental data is obtained for a parameterization of the cavity that is fitted[47] to reproduce an interface-related observable, i.e. the theoretical estimate of the PZC of Pt(111). The original parameterization of the SSCS cavity has been found instead to produce a relatively good agreement with experimental data without the need of refitting.

While it is important to stress how the different approaches can be extended and tuned to improve the description of electrochemical systems, it is worth pointing out that the reported analysis is based on continuum models that only account for part of the physical phenomena occurring at electrified interfaces. At the center of the reported analysis is the description of the diffuse layer, of its shape and characteristics. Nonetheless, a more realistic model should account for the different sizes of the ions composing the electrolyte. Moreover, as it is also clear from the results reported, the dielectric properties of the liquid

solution at the interface with a solid substrate need to be properly modeled in order for the continuum approach to be meaningful. The bare substrate and, even more, a charged interface will induce order and rigidity in the overlaying liquid, substantially affecting the dielectric permittivity over a distance of one or more solvation layers. The fact that current state-of-the-art continuum models are not able to describe with the same accuracy systems with different charge states, and in particular require a separate parameterization for anions, is clearly a limitation of the current techniques in dealing with electrified interfaces. Similarly, non-linear effects in the dielectric response of the liquid may account for some of the deviations observed at higher applied potentials. Other possibly minor effects that are not explicitly accounted for in the presented models are the ones related to the change in dielectric screening of the electrolyte solution for high concentrations of the diffuse layer. Cancellation of errors resulting from the parameterization of the model may lead to an approach that seems accurate, but lacks transferability. As more and more ingredients are added and carefully tuned, they will unlock the full potential of continuum models for electrochemical setups.

#### 4.5. Kohn-Sham Potential and Force Contributions

In order to perform fully self-consistent calculations with the numerical electrolyte models discussed in the paper, one needs to add contributions that arise from the dependence of the free energy on the solute electron density  $\rho^{el}(\mathbf{r})$  and on the atomic positions  $\{\mathbf{R}_\alpha\}$  to the Kohn-Sham (KS) potential  $V_{KS}(\mathbf{r})$  and to the atomic forces  $\{\mathbf{f}_\alpha\}$ , respectively. This procedure leads to contributions analogous to the ones that appear in a standard vacuum environment, but with additional terms that arise from the dependence of the free-energy functional on the interface function  $s(\mathbf{r})$ . These interface-related terms can be easily computed exploiting the availability of the functional derivative of the free energy  $F[s(\mathbf{r})]$  with respect to  $s(\mathbf{r})$ :

$$(4.45) \quad V_{KS}^s(\mathbf{r}) = \int \frac{\delta s(\mathbf{r}')}{\delta \rho^{el}(\mathbf{r})} \frac{\delta F[s(\mathbf{r})]}{\delta s(\mathbf{r}')} d\mathbf{r}',$$

$$(4.46) \quad \mathbf{f}_\alpha^s = \int \frac{\partial s(\mathbf{r}')}{\partial \mathbf{R}_\alpha} \frac{\delta F[s(\mathbf{r})]}{\delta s(\mathbf{r}')} d\mathbf{r}'.$$

For the Poisson-Boltzmann (PB) model, its linear-regime approximation (linearized-PB, LPB) and its size-modified variant (size-modified PB, MPB), the functional derivatives  $\delta F[s(\mathbf{r})]/\delta s(\mathbf{r})$  take the following forms (see Eqs. 9, 15 and 22 in the main paper for the corresponding free-energy expressions):

$$(4.47) \quad \frac{\delta F^{PB}[s(\mathbf{r})]}{\delta s(\mathbf{r})} = (\varepsilon_0 - 1) \frac{|\nabla\phi(\mathbf{r})|^2}{8\pi} + k_B T \sum_{i=1}^p c_i^0 e^{-\frac{z_i\phi(\mathbf{r})}{k_B T}},$$

$$(4.48) \quad \frac{\delta F^{LPB}[s(\mathbf{r})]}{\delta s(\mathbf{r})} = (\varepsilon_0 - 1) \frac{|\nabla\phi(\mathbf{r})|^2}{8\pi} + \frac{\sum_{i=1}^p z_i^2 c_i^0}{2k_B T} \phi^2(\mathbf{r}) + k_B T \sum_{i=1}^p c_i^0,$$

$$(4.49) \quad \frac{\delta F^{MPB}[s(\mathbf{r})]}{\delta s(\mathbf{r})} = (\varepsilon_0 - 1) \frac{|\nabla\phi(\mathbf{r})|^2}{8\pi} - k_B T c_{max} \ln \left( c_{max} - \sum_{i=1}^p c_i^0 \right) + \\ + k_B T c_{max} \ln \left( c_{max} - \sum_{i=1}^p c_i^0 \left( 1 - e^{-\frac{z_i\phi(\mathbf{r})}{k_B T}} \right) \right).$$

It is evident from Eqs. 4.45 and 4.46 that the interface-related contributions to the KS-potential and to the atomic forces vanish if the interface function does not explicitly depend on  $\rho^{el}(\mathbf{r})$  and  $\{\mathbf{R}_\alpha\}$ , respectively. Therefore, the  $\{\mathbf{f}_\alpha^s\}$  are zero if the density-based cavity function from the self-consistent continuum solvation (SCCS) model is employed,  $V_{KS}^s(\mathbf{r})$  is zero if the solute cavity is constructed from atom-centered spheres as in the soft-sphere continuum solvation (SSCS) model, and both  $\{\mathbf{f}_\alpha^s\}$  and  $V_{KS}^s(\mathbf{r})$  are zero for the space-fixed planar interface function.

#### 4.6. Analytical solution of the Poisson-Boltzmann model of the diffuse layer

Starting from the Poisson-Boltzmann equation for a two-dimensional system homogeneous in the plane,

$$(4.50) \quad \frac{d^2\phi(x)}{dx^2} = -\frac{4\pi}{\varepsilon_0} \left( \rho(x) + \sum_{i=1}^p z_i c_i(\phi(x)) \right)$$

one can apply the chain rule to the left hand side:

$$(4.51) \quad \frac{d^2\phi}{dx^2} = \frac{d}{d\phi} \frac{d\phi}{dx} \frac{d\phi}{dx} = \frac{1}{2} \frac{d}{d\phi} \left( \frac{d\phi}{dx} \right)^2.$$

The right hand side can be rewritten using the result from the variational approach applied to the ionic concentrations

$$(4.52) \quad c_i(x) = c_i^0 e^{-\frac{z_i \phi(x)}{k_B T}}$$

In the region of interest,  $\rho(x)$  is taken to be zero, and the electrolyte is made up of two species of equal and opposite charges, and equal concentrations. Hence the right hand side of the equation can be expressed as

$$(4.53) \quad -\frac{4\pi z c^0}{\epsilon_0} \left( e^{-\frac{z\phi}{k_B T}} - e^{\frac{z\phi}{k_B T}} \right),$$

or in terms of hyperbolic functions,

$$(4.54) \quad \frac{8\pi z c^0}{\epsilon_0} \sinh\left(\frac{z\phi}{k_B T}\right).$$

Performing integrations on both sides with respect to  $\phi$  gives the following differential equation

$$(4.55) \quad \left(\frac{d\phi}{dx}\right)^2 = \frac{8z c^0 \pi}{\epsilon_0} \sinh\left(\frac{z\phi}{k_B T}\right)$$

$$(4.56) \quad = \frac{16\pi k_B T c^0}{\epsilon_0} \cosh\left(\frac{z\phi}{k_B T}\right) + k$$

where  $k$  is the constant of integration, which can be immediately found by imposing the condition  $(d\phi/dx)(|x| \rightarrow \infty) = \phi(|x| \rightarrow \infty) = 0$ . In this form, the first-order differential equation can be simplified by a double-angle identity,

$$(4.57) \quad \frac{d\phi}{dx} = -\left(\frac{32\pi k_B T c^0}{\epsilon_0}\right)^{\frac{1}{2}} \sinh\frac{|z|\phi}{2k_B T},$$

where the negative sign is chosen to match the physical interpretation of the electric field.

Rearranging terms and integrating gives:

$$(4.58) \quad \ln\left(\kappa \left| \tanh\left(\frac{|z|\phi}{4k_B T}\right) \right| \right) = -\left(\frac{32\pi k_B T c^0}{\epsilon_0}\right)^{\frac{1}{2}} x,$$

where the constant of integration  $\kappa$  has been moved into the logarithm for convenience. Note that  $x$  is limited to positive values but since the setup is inherently symmetric, it should follow that the expression for potential is also symmetric; hence the absolute value of  $x$

should be taken. The expression from this point is inverted and the potential can finally be written as

$$(4.59) \quad \phi(x) = \frac{4k_B T}{|z|} \coth^{-1} \left( \kappa e^{\frac{32\pi k_B T c^0}{\epsilon_0} |x|} \right),$$

where  $\kappa$  can be found by imposing conditions at the electrolyte interface.

#### 4.7. Analytical solution of the linearized Poisson-Boltzmann model of the diffuse layer

Searching the solution to the linearized Poisson-Boltzmann equation

$$(4.60) \quad \nabla \cdot \varepsilon(\mathbf{r}) \nabla \phi(\mathbf{r}) - k^2 \gamma(\mathbf{r}) \phi(\mathbf{r}) = -4\pi \rho(\mathbf{r}),$$

in the region where  $\epsilon(\mathbf{r}) = \epsilon_0$ ,  $\gamma(\mathbf{r}) = 1$ , and  $\rho(\mathbf{r}) = 0$ , for a two-dimensional system that is homogeneous in the plane, corresponds to solving the following one-dimensional differential equation

$$(4.61) \quad \frac{d^2 \phi(x)}{dx^2} - \frac{k^2}{\epsilon_0} \phi(x) = 0.$$

This is a trivial differential equation with the general solution

$$(4.62) \quad \phi(x) = c_1 e^{\frac{kx}{\sqrt{\epsilon_0}}} + c_2 e^{-\frac{kx}{\sqrt{\epsilon_0}}},$$

where  $c_1$  and  $c_2$  are constants determined by appropriate boundary conditions. Applying the condition that  $\phi(|x| \rightarrow \infty) = 0$ ,

$$(4.63) \quad \phi(x) = c e^{-\frac{k|x|}{\sqrt{\epsilon_0}}},$$

where  $c$  is determined by the electric field at the electrolyte interface.

## CHAPTER 5

### FIELD-AWARE INTERFACES IN CONTINUUM SOLVATION

1

#### 5.1. Introduction

Recent efforts to account for charged solutes have benefited from a similar non-local description for the interface. Previously discussed was the ability for the local interface models to successfully describe charged systems with a simple reparameterization. The disadvantage of this method is in the reparameterization itself, which is costly and unsuitable if investigating variable charge systems.

There are a number of strategies for approaching this problem, of which not all exploit a non-local interface. One such idea is to represent some of the solvent explicitly instead of fully implicitly. These approaches have a lot of potential, however, an automatic decision of what solvent molecules should be represented is non-trivial, and the configuration of this single layer of water needs to be considered. Therefore, there is merit in developing an implicit solution to the problem, and this is supported by recent additions to the literature.

A more hands-off approach is to infer the charge/polarization state of the embedded QM system from its properties instead of using chemical intuition. It can be argued that the electrostatic field generated by polarization or global charge within a solute can serve as an effective proxy for this charged state. Multiple groups have pursued this idea. Pomogeva and Chipman proposed a field-extremum model that calculates the electric field normal to the interface boundary and selects the minimum and maximum values as potential hydrogen bond locations. This would be added to the free-energy outside the simulation itself as a post processing term. The rationale for using the normal component of the electric field is Gauss' law. For a closed surface with enclosed charge, the electric flux through the surface is proportional to the charge. Therefore by measuring the electric field normal to the surface,

---

<sup>1</sup>The majority of this chapter is reproduced from Matthew Truscott and Oliviero Andreussi. Field-aware interfaces in continuum solvation. *J. Phys. Chem. B* 2019, 1213, 16, 3513-3524. DOI: <https://doi.org/10.1021/acs.jpcc.9b01363>, with permission from the American Chemical Society

a measure of the charge is obtained, and a measure of the effective charge distribution on the surface of the interface is obtained. This not only means we have the total charge of the system, but also the distribution of that charge. Pomogeva and Chipman used this knowledge to infer likely positions for hydrogen bonds, or in other words, locations on the interface where interactions between solvent and solute were closest.

Rahimi and collaborators took a similar approach, again treating the electric field normal to the interface as part of a correction factor that would be added as a post processing term. Sundararaman and Goddard proposed a more ambitious approach, where the interface adjusted as a function of the normal electric field. A changing interface during the electronic density optimization introduces more potential for instability but results in the desired self-adjustment for localized charge that is desirable in continuum solvation.

Following from these successes, the research members of the Materialab group, primarily myself and Dr. Andreussi, set out to extend the capabilities of the SCCS and soft-sphere models with this modification derived from the electric field. The modification would act on the interface in a self-consistent fashion. The electronic minimization algorithm implemented for these models is an iterative one, and the new interface function will now be a function of the electronic density as part of the iterative process. This cyclic dependency is rationalized by the stable nature of the interface function.

Of the two models, the soft-sphere approach has a more straightforward way to account for the electric field. We consider the electric field normal to each of the exposed spheres in the molecule, and adjust the sizes of each atomic-centered sphere accordingly.

The vacuum electric field normal to the interface can be expressed by

$$(5.1) \quad E^n(\mathbf{r}) = \mathbf{E}(\mathbf{r}) \cdot \nabla s(\mathbf{r})$$

where the electric field in vacuum  $\mathbf{E}(\mathbf{r})$  can be computed from the solute charge density, summing electrostatic and ionic density terms,



$$(5.2) \quad \mathbf{E}(\mathbf{r}) = \int \mathbf{G}(\mathbf{r} - \mathbf{r}') \left( \rho^{\text{el}}(\mathbf{r}) + \sum_a z_a \delta(|\mathbf{r} - \mathbf{R}_a|) \right) d\mathbf{r}'$$

where  $\mathbf{G}(\mathbf{r} - \mathbf{r}')$  depends on the boundary conditions of the problem. The unscaled soft-sphere function is defined as before, eq. 2.12, and denoted by  $h^0(r)$ . In order to work with the electric field expression above, we must first derive the gradient of the interface function,

$$(5.3) \quad \nabla s(\mathbf{r}) = \nabla \left[ 1 - \prod_a h_a^0(r_a) \right] = - \sum_a \nabla h_a^0(r_a) \prod_{b \neq a} h_b^0(r_b) = (s(\mathbf{r}) - 1) \sum_a \frac{\nabla h_a^0(r_a)}{h_a^0(r_a)}$$

where

$$(5.4) \quad \begin{aligned} \nabla h_a^0(r_a) &= \frac{dh_a^0}{dr_a} \nabla r_a \\ &= \frac{1}{\sqrt{\pi} \Delta} \exp \left( - \frac{(r_a - \alpha_\xi R_a^{\text{vdW}})^2}{\Delta_\xi^2} \right) \frac{\mathbf{r}_a}{r_a} \end{aligned}$$

The electric flux through each atomic centered sphere can therefore be defined,

$$(5.5) \quad \Phi_a = \int E_a^n(\mathbf{r}) d\mathbf{r} = - \int \mathbf{E}(\mathbf{r}) \cdot \nabla h_a^0(r_a) \prod_{b \neq a} h_b^0(r_b) d\mathbf{r}$$

We now have a scalar quantity that describes the electric flux through each atom centered sphere. For atoms away from the surface, the flux is zero, since the electric field normal to the surfaces here is zero. Therefore, only surface exposed spheres should shrink in size. A new interface function can be defined as a function of this flux term, and the result is an interface function that relies on both the electronic and ionic components of the system. Choosing the electric flux through each atom has a good physically intuitive argument, Gauss' law states that the total flux passing through a closed surface is proportional to the enclosed charge.

A definition of a model that depends on the flux differs somewhat from predecessor models, for example, the CANDLE model where the scaling is a direct function of the normal

electric field, however it is the natural extension of such ideas when applied to the soft-sphere interface function. The product of soft-sphere functions on the right hand side of equation 5.5 ensures that the flux is not computed for the portions of the sphere contained inside neighbouring spheres. This choice makes the connection between the flux and the contained charge less direct, especially for soft spheres that have only a small fraction of their surface exposed to the continuum. This does however mirror the expected lack of effect that less exposed atoms should have on the interface. Indeed, atoms that lie buried within molecules would not be expected to have significant interaction with the solvent, since their electronic density is masked by surrounding solute atoms.

New definitions of the interface and soft-sphere functions are similar to before (compare with equations 2.11 and 2.12), only now with a multiplicative scaling factor,

$$(5.6) \quad \hat{s}(\mathbf{r}, \mathbf{R}_i) = 1 - \left[ \prod_i h_i^{\Phi_i}(\{\xi\}; \|\mathbf{r} - \mathbf{R}_i\|) \right]$$

and

$$(5.7) \quad h_i^{\Phi_i}(\alpha_\xi, R_i^{\text{vdW}}, \Delta_\xi; \|\mathbf{r} - \mathbf{R}_i\|) = \frac{1}{2} \left[ 1 + \operatorname{erf} \left( \frac{\|\mathbf{r} - \mathbf{R}_i\| - f(\Phi_i)\alpha_\xi R_i^{\text{vdW}}}{\Delta_\xi} \right) \right]$$

Note that with this definition of the interface function, there exists a cyclic dependence between the fluxes and the interface function. This can be resolved by a self-consistent loop, but it would add unnecessary computational complexity to the problem. Instead, the unscaled soft spheres  $h_i^0$  are adopted for the calculation of the fluxes. This should produce a result identical to any reasonable margin of error: the electronic density does not dramatically vary in the vicinity of the sphere surfaces and therefore the change in flux by modifying the size of the sphere by some incremental amount should be negligible.

## 5.2. Functional Derivatives

In order to be able to exploit the new interface function in the Kohn-Sham SCF procedure and for geometry optimization, we need to characterize the derivatives of the

field-aware soft-sphere interface with respect to the electronic density and ionic positions, since the field fluxes depend nonlocally on both the electronic and ionic charge densities.

We can express the functional derivative of the interface function with respect to the electronic density in terms of the individual field fluxes and their functional derivatives,

$$(5.8) \quad \frac{\delta \hat{s}}{\delta \rho^{el}}(\mathbf{r}, \mathbf{r}') = \sum_a \frac{\partial \hat{s}}{\partial \Phi_a}(\mathbf{r}') \frac{\delta \Phi_a}{\delta \rho^{el}}(\mathbf{r}).$$

By applying this result to Eq. (2.7), we get for the interface contribution to the Kohn-Sham potential

$$(5.9) \quad \begin{aligned} V_{\text{KS}}^{interface}(\mathbf{r}) &= \int \sum_a \frac{\partial \hat{s}}{\partial \Phi_a}(\mathbf{r}') \frac{\delta \Phi_a}{\delta \rho^{el}}(\mathbf{r}) \frac{\delta E^{interface}[s]}{\delta s}(\mathbf{r}') d\mathbf{r}' \\ &= \sum_a \frac{\delta \Phi_a}{\delta \rho^{el}}(\mathbf{r}) \int \frac{\partial \hat{s}}{\partial \Phi_a}(\mathbf{r}') \frac{\delta E^{interface}[s]}{\delta s}(\mathbf{r}') d\mathbf{r}' \end{aligned}$$

Note that although the functional derivative of the interface is defined in  $\mathbf{r}$  for a change in the charge distribution given in  $\mathbf{r}'$ , these coordinate systems can be separated cleanly. This is particularly useful for implementation, since it would be numerically cumbersome and expensive to manage a quantity defined on six dimensions. By separating the expression as we have shown, one can now evaluate a three-dimensional integral and multiply by the final functional derivative term. The partial derivative of the interface with respect to the field fluxes can be decomposed as

$$(5.10) \quad \begin{aligned} \frac{\partial \hat{s}}{\partial \Phi_a}(\mathbf{r}') &= \frac{\partial \hat{s}}{\partial h_a^{\Phi_a}} \frac{dh_a^{\Phi_a}}{df}(r'_a) \frac{df}{d\Phi_a} \\ &= \alpha_\xi R_a^{\text{vdW}} \frac{dh_a^{\Phi_a}}{dr'_a} \frac{df}{d\Phi_a} \prod_{b \neq a} h_b^{\Phi_b}(r'_b) \end{aligned}$$

where the total derivative of  $h_a$  is given in Eq. (5.4), while the derivative of the field-aware function  $df/d\Phi_a$  will depend on its definition (as reported in the following Section).

The partial derivatives of the interface function with respect to the ionic positions is given by

$$\begin{aligned}
(5.11) \quad \nabla_{\mathbf{R}_c} \hat{s}(\mathbf{r}) &= - \sum_a \nabla_{\mathbf{R}_c} h_a^{\Phi_a}(r_a) \prod_{b \neq a} h_b^{\Phi_b}(r_b) \\
&= \sum_a \frac{1}{\sqrt{\pi} \Delta} \frac{dh_a}{dr_a} \left( \delta_{ca} \frac{\mathbf{r}_a}{r_a} - \alpha R_a^{\text{vdW}} \frac{df}{d\Phi_a} \nabla_{\mathbf{r}_c} \Phi_a \right) \prod_{b \neq a} h_b^{\Phi_b}(r_b) \\
&= \nabla_{\mathbf{R}_c} s(\mathbf{r}) + \sum_a \nabla_{\mathbf{R}_c} \Phi_a \frac{\partial \hat{s}}{\partial \Phi_a}(\mathbf{r})
\end{aligned}$$

Here, the result can be expressed as a correction to the derivative of the original SSCS interface function. When combined with Eq. (2.8), the field-aware contribution to interatomic forces can be written as

$$(5.12) \quad \mathbf{f}_c^{\text{interface}} = \mathbf{f}_c^{\text{interface},0} - \sum_a \int \nabla_{\mathbf{R}_c} \Phi_a \frac{\partial \hat{s}}{\partial \Phi_a}(\mathbf{r}) \frac{\delta E^{\text{interface}}[s]}{\delta s}(\mathbf{r}) d\mathbf{r}$$

Comparing Eq.s (5.12) and (5.9) allows to appreciate the similarities between the two expressions.

The last ingredients missing in the above expressions are the derivatives of the field fluxes with respect to the solute degrees of freedom. In particular, the functional derivative of the field fluxes with respect to the electronic density can be expressed as

$$\begin{aligned}
(5.13) \quad \frac{\delta \Phi_a}{\delta \rho^{\text{el}}}(\mathbf{r}) &= \frac{\delta}{\delta \rho^{\text{el}}} \int \mathbf{E}(\mathbf{r}') \cdot \nabla h_a^0(\mathbf{r}') \prod_{b \neq a} h_b^0(r'_b) d\mathbf{r}' \\
&= \frac{\delta}{\delta \rho^{\text{el}}} \int \left[ \int \mathbf{G}(\mathbf{r}' - \mathbf{r}'') \left( \rho^{\text{el}}(\mathbf{r}'') + \sum_c z_c \delta(r''_c) \right) d\mathbf{r}'' \right] \cdot \nabla h_a^0(\mathbf{r}') \prod_{b \neq a} h_b^0(r'_b) d\mathbf{r}' \\
&= \int \mathbf{G}(\mathbf{r}' - \mathbf{r}) \cdot \nabla h_a^0(\mathbf{r}') \prod_{b \neq a} h_b^0(r'_b) d\mathbf{r}'
\end{aligned}$$

It is important to note that the final result above is free from any dependence on the electronic density. This allows us to compute this factor only once, before the SCF loop, for a given set of atomic positions. Obtaining the derivative of the flux with respect to the ionic positions is more long-winded and so we leave the steps to the supporting information, showing only the initial expansion and the final result

$$\begin{aligned}
(5.14) \quad \nabla_{\mathbf{R}_c} \Phi_a &= \nabla_{\mathbf{R}_c} \int \mathbf{E}(\mathbf{r}) \cdot \nabla h_a(r_a) \prod_{b \neq a} h_b(r_b) d\mathbf{r} \\
&= - \int (\mathbf{H}[\phi(z_c \delta(r_c))] \nabla h_a(r_a) + \delta_{ac} \mathbf{H}[h_a(r_a)] \mathbf{E}(\mathbf{r})) \prod_{b \neq a, c} h_b(r_b) d\mathbf{r} \\
&\quad + (\delta_{ac} - 1) \int [\mathbf{E}(\mathbf{r}) \cdot h_a(r_a)] \nabla h_c(r_c) \prod_{b \neq a, c} h_b(r_b) d\mathbf{r}
\end{aligned}$$

expressed in terms of the Hessian  $\mathbf{H}[\phi(z_c \delta(r_c))]$  of the electrostatic potential of atom  $c$  and of the Hessian  $\mathbf{H}[h_a(r_a)]$  of the soft sphere on atom  $a$ .

All of the derivatives reported above were tested by comparing their analytic expressions with numerical finite-differences estimates.

### 5.3. Field-aware function

The main ingredient of the field-aware approach is represented by the scaling function that is responsible to shrink the solvation radii as the field flux increases. This function should assume a value of 1 for low positive or negative electric fluxes, so as to not rescale the continuum interface for non-polar atoms or non-charged systems. Alternatively, in order to reproduce the shrinking of the interface close to polar/charged residues, the field-aware function needs to decrease for high positive or negative electric fluxes. Moreover, as demonstrated in multiple preceding publications (e.g. in Ref. ), the mean average error from anions is greater than cations, suggesting an asymmetric relationship between the flux and the size of the atom in question. For the above reasons, we chose to introduce a field-aware function using a piece-wise definition similar to the one adopted for the SCCS model, namely

$$(5.15) \quad f(\Phi) = 1 - f_0 t(\Phi) (\kappa - \text{sgn}(\Phi))^2$$

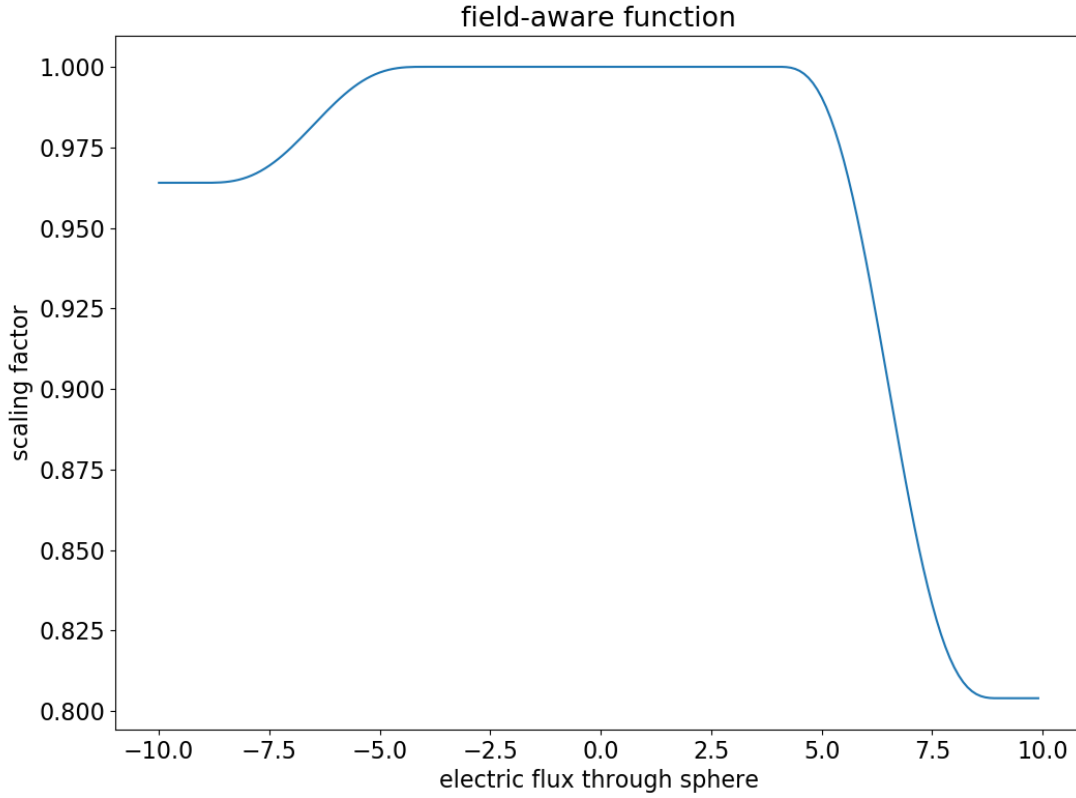


FIGURE 5.1. The field aware function with parameters set to demonstrate the asymmetry with a non-zero field factor. This smoothly varying function can be tuned to ignore atoms with insufficient flux, which can improve the stability of the self-consistent electrostatic calculation.

where  $t(\Phi)$  is a piece-wise function defined by

$$(5.16) \quad t(\Phi) = \begin{cases} 0 & \Phi \leq \Phi_{\min} \\ \frac{1}{2\pi}(x - \sin x) & \Phi_{\min} < \Phi < \Phi_{\max} \\ 1 & \Phi \geq \Phi_{\max} \end{cases}$$

$$(5.17) \quad x = 2\pi \frac{\Phi - \Phi_{\min}}{\Phi_{\max} - \Phi_{\min}}$$

There are four field-specific parameters to be tuned. The maximum scaling of soft

spheres corresponding to positive and negative residues is controlled by the combination of a global field-aware intensity,  $f_0$ , and an asymmetry parameter,  $\kappa$ . The transition region between a non field-aware regime and the maximum scaling factor is controlled by two thresholds,  $\Phi_{\min}$  and  $\Phi_{\max}$ , of the electrostatic fluxes. A typical example of the behavior of the field-aware function is reported in Figure 5.1.

In principle, since this function directly affects the scaling of the soft-spheres, which was previously globally tuned for all species of a specific charge, a retuning of the SSCS  $\alpha_\xi$  scaling factor is also necessary. In the subsequent analysis, we find it convenient to recast the field cutoff parameters in terms of a mean flux, and a spread parameter. These are related by  $\Phi_{\text{avg}} = (\Phi_{\max} - \Phi_{\min})/2$  and  $\Phi_{\text{spread}} = \Phi_{\max} - \Phi_{\text{avg}}$ .

It should also be noted that the flux values are given in Atomic Rydberg units.

#### 5.4. Computational Details

In order to validate this newly defined field-aware model, a parameterization effort is necessary. Solvation energy is calculated for a training set of small molecules, and an attempt is made to minimize the mean absolute error of this energy with respect to experimental values.

Pseudopotentials were chosen according to the SSSP (standard solid state pseudopotentials) efficiency set. Kohn-Sham wavefunctions and charge densities were expanded in plane waves up to kinetic energy cutoffs of 30 and 300 Ry respectively. Simulation cells were chosen to be 30.0 a.u. to reduce periodic boundary artefacts in charged systems. Note that when displaying quantities of flux, these are given in Atomic Rydberg units unless specified. Hence if a sphere has a flux of  $8\pi$ , the corresponding charge housed within is equivalent to one electron charge. A few selected neutral and charged molecules were studied as a function of the field-aware parameters and the soft-sphere scaling factor. The studied systems were chosen so as to be representative of the different pathological cases encountered in standard continuum models. In particular, water and its corresponding ionic species (hydroxyl anion and hydronium cation) were chosen due to the fact that the solvation free energies of both charged species show large MAE values with the soft-spheres model, even

after re-parameterization for charged molecules was performed. In addition, one neutral molecule, one cation, and one anion are also picked to show the potential of this model. These are chosen from a set of 13 neutral molecules, 15 cations and 15 anions, as the worst performing species when calculating the solvation energy by the SSCS. Following this criterion, propanoic acid, the diethylether cation, and the benzyl alcohol anion were selected. Implementation and simulations were performed using a development version of the Environ-1.0 plugin for multiscale embeddings in condensed-matter simulations, as coupled with the Quantum Espresso (v6.3) package.

As discussed previously, the field aware function has four tunable parameters. In addition, the soft-sphere parameters, which for the neutral set were chosen as a best fit including species with internal polarization, need to be reconsidered. For simplicity we begin by just considering how the scaling factor changes. Hence, for preliminary analysis of neutral molecules, we consider values of  $\alpha_\epsilon$  ranging from 1.12 to 1.20. This is under the assumption that the optimal scaling factor from the soft-spheres is a result of a compromise between molecules without internal polarization effects, which should thus require a larger spherical functions, and molecules with internal polarization effects, which should require smaller spherical functions for specific atoms.

For cations and anions, the asymmetry is set to -1 and 1, respectively, to minimize the computational costs of gathering results, since the flux values for the species we test for all have the same sign, negative for anions and positive for cations. A range of electric flux parameters compatible with the values observed for the selected molecules was considered (see fluxes reported in Figure 5.2), again to minimize computational time. This range may need to be extended for a more comprehensive parameterization of the model that includes multivalence charged systems.

The field awareness  $f_0$  is chosen to be the upper limit on the reduction of the radii of the spherical functions. This was originally chosen to match the reduction as seen in the soft-spheres model between the neutral set and the anion set, but considering the reparameterization of this scaling as part of the field-aware model, it was found that increasing this






FIGURE 5.2. Comparison of electric flux values for each individual atom for the six molecules. A color map proportional to the absolute value of the flux is adopted. Atoms showing the largest absolute electric fluxes are identified in Figure 5.3.

Molecule	Abs Error	Asymmetry	$\alpha$	$f_0$	Flux Mean	Flux Spread
Water	0.0163	0.0	1.16	0.06	1.5	0.6
Propanoic Acid	0.1035	1.0	1.20	0.05	2.0	0.4
Water (+)	0.2950	1.0	1.18	0.05	1.5	0.6
Diethyl Ether (+)	1.2634	1.0	1.12	0.06	1.5	0.6
Water (-)	0.4813	-1.0	1.20	0.07	1.5	0.6
Benzyl Alcohol (-)	10.9610	-1.0	1.12	0.07	2.0	1.5

TABLE 5.1. Best Parameter Set

value yielded better results. Note that as defined, both the asymmetry factor and the field awareness control the global intensity of the field-aware procedure.

## 5.5. Results

The field-aware model is able to converge for the majority of parameter combinations, although numerical issues arise with the SCF procedure when specific ranges of parameter values are selected. This is linked to the sharpness of the field-aware potential when the field-aware function has a transition that is too sharp with respect to the flux. Nevertheless, for the ranges of parameters tested in this study, involving about 3000 different combinations,




FIGURE 5.3. Set of molecules tested. The atom possessing the maximum absolute electric flux is highlighted by a red circle, and its flux value is reported

99% of the SCF and geometry optimization calculations converged successfully, validating the analytic differentiation of field-aware interfaces reported above.

Due to the number of parameters, it is challenging to optimize this model from scratch. Our main aim here is thus to display insights from a broad sweep, as well as observations that may simplify the parameterization process. As this process will benefit from a significant number of experimental and theoretical results on solvation of charged molecules and interfaces (e.g. results of differential capacitance of metal substrates), its completion lies outside of the scope of this work.

We start by showing in Figure 5.2 the final values of the electric flux through each soft sphere for the selected molecules. For full numerical values, we refer to the Supporting Information. The values attained for the molecules chosen match expectations, neutral

molecules center around zero flux, with the oxygen atom being where most of the negative charge imbalance lies and anions show large relative charge differences. One can see how the field-aware function acts on individual atoms by comparing Figure 5.2 to the field-aware function, as in Figure 5.1. Any flux with an absolute value over the cutoff will result in a minimum value for the scaling factor, which varies from 1 (no scaling) to a value determined by the field awareness and asymmetry parameters. In particular, due to the definition of the asymmetry factor, a deviation of 0.5 away from symmetry produces a heavy bias towards spheres of negative flux. The choice of flux cutoff parameters can be simplified by considering typical flux values for a training set. For this investigation we limit our analysis to consider how an adjustment of parameters for individual species can improve on absolute solvation error values. In Table 5.1 we report the best performing set of parameters for each considered compound. Despite most of these parameters fall in a similar range of values, the corresponding field-aware functions are clearly not representative of the values of a fully optimized scaling function. For example, in the case of the chosen anions, the absolute flux through the oxygen atoms is much greater than the other atoms and as such, only the sphere around this atom is appreciably altered.

The kind of parameterization required will also depend on the broadness of the task for which the field-aware function is fitted. A conservative approach to the field-aware function would be to tune it so that it matches the shrinking of the interfaces observed for charged species. In particular, Fisticaro et al. reported a optimized value of the homogeneous soft-sphere scaling parameter  $\alpha_\xi$  that goes from a value of 1.12 for neutral compounds to a value of 1.1 for cations, reaching 0.98 for anions. To reproduce this trends, a field-aware function with sharp transitions between low and high flux values would be sufficient. However, if the field awareness is designed to also improve the description of polar groups in neutral molecules, a more smoothly varying field-aware function needs to be considered.

#### 5.5.1. Propanoic Acid

Propanoic acid (figure 5.4) is the worst performing neutral molecule for the standard SSCS model, out of the 13 testing neutral molecules. This acidic compound is expected

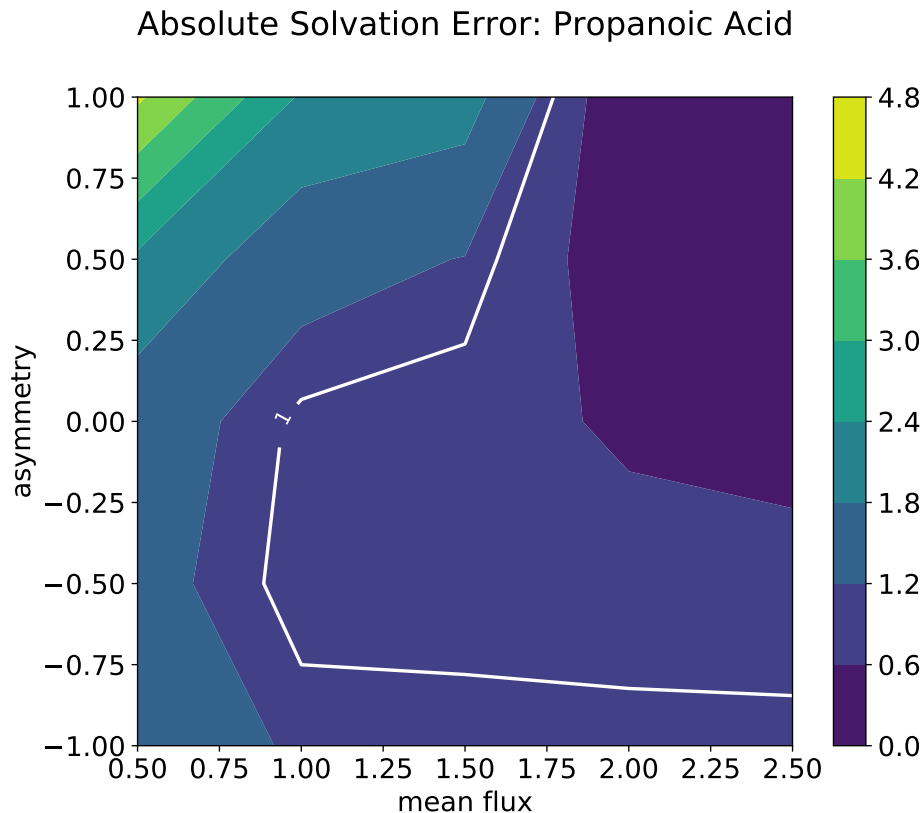


FIGURE 5.4. Absolute solvation free energy error for propanoic acid in kcal/mol as a function of the mean flux threshold and the asymmetry parameter. The flux spread is fixed at 0.4, the field awareness is fixed at 0.05, while an homogenous scaling parameter  $\alpha_\xi = 1.20$  is considered

to produce larger errors, due to its polar nature. The SSCS model underestimates the solvation energy by 2.53kcal/mol, suggesting a scaling factor that is too small. Having a function that reduces the size of the spherical functions would not be able to improve on the SSCS result and would only worsen the result. The problem is that the homogeneous scaling factor optimized in the original SSCS model represents a trade-off between apolar and polar groups, and is thus on average smaller the needed for apolar groups, while being larger than needed for polar regions. As a result, when the global scaling factor is allowed to vary above the original soft-sphere value, improved results are observed. In particular, in the region we considered in our five dimensional parameter search, the global scaling factor




FIGURE 5.5. Propanoic Acid. Unscaled interface (left) represents the spherical functions as calculated by standard SSCS. Scaled interface (right) is based off an optimized set of parameters, as reported in Table 5.1.

and the asymmetry had the greatest influence on the solvation energy for this acid. In fact, a reparameterization with  $\alpha = 1.20$  brought the calculated solvation energy remarkably close to the experimental value (see results in Table 5.1), a result that is consistent with the effect of interface size on the solvation energy. In addition, promising trends are shown in Figure 5.5, featuring the behavior of the absolute error on solvation free energies against the asymmetry and the average flux values. This figure shows that for the chosen fixed parameters, putting less weight on negative fluxes provides particularly small errors. Indeed, the region where the asymmetry is negative appears to be more stable with respect to the flux average parameter. This is consistent with the fact that there is a single oxygen atom with appreciable negative flux, which is always accounted for in the range chosen for the maximum cutoff flux.

### 5.5.2. Diethyl Ether

As seen in Figure 5.2, this cation has little variance in flux among its atomic constituents. In the standard SSCS model, the solvation free energy of diethyl ether is overestimated by 16.57kcal/mol even after scaling. Despite the good performance of the SSCS model after reparameterization for the cations, there remain a number of species that have significant solvation errors. In Figure 5.6 we plot the field awareness against the flux average

### Absolute Solvation Error: Diethyl Ether Cation




FIGURE 5.6. Absolute solvation free energy error for diethyl ether in kcal/mol, as a function of mean flux threshold and field awareness parameter. The flux spread is fixed at 0.6, and an homogenous scaling factor of  $\alpha_\xi = 1.20$  is adopted. An asymmetry  $\kappa = 1$  is selected, as in this molecule there are no atoms with negative fluxes.

to show how the absolute error for this species can be drastically improved without a modification of  $\alpha_\xi$  to a lower value, which is the standard strategy when accounting for cations. The region to the right is one where the flux values are too low to be considered by the field-aware function and thus the error is equal to the standard SSCS result for  $\alpha_\xi = 1.12$ . Note that since the minimum and maximum possible scaling is controlled by the choice of parameters under investigation, as long as the model is tuned sensibly, the worst case scenario ought to be as one can see in the region to the right of this plot, i.e. no improvement. On the opposite end, when parameters such as the ones varied in the plot are properly chosen, the absolute solvation error can improve from 18.07kcal/mol to 1.35kcal/mol. Despite the




FIGURE 5.7. Diethyl Ether. Unscaled interface (visualized on the left half of the molecule) represents the spherical functions as calculated by standard SSCS. Scaled interface (represented on the right half of the molecule) is based off an optimized set of parameters, as reported in Table 5.1. The effects of the field-aware scaling are small, mostly noticeable for the acidic hydrogen atom bonded to the oxygen.

large change in solvation free energy, the overall effect on the shape of the interface is mostly noticeable for the acidic hydrogen atom bonded to the oxygen, as reported in Figure 5.7.

### 5.5.3. Benzyl Alcohol

For the range of parameters we currently have tested for, a potential optimal value for this particular species is somewhat far off. Nonetheless, like in the diethyl ether case, a significant improvement of the results can be obtained by properly selecting the field-

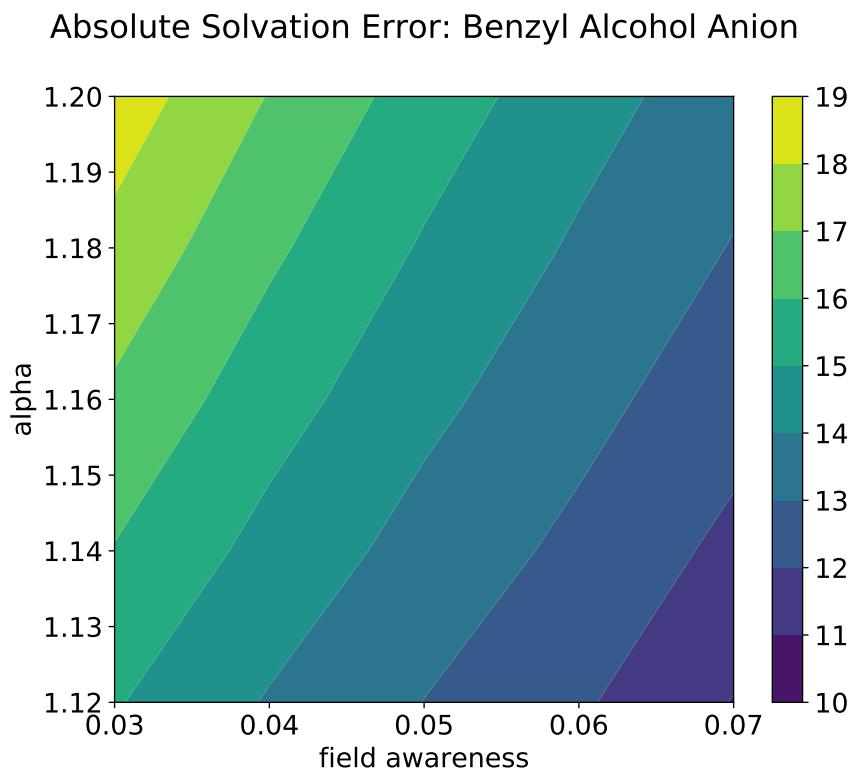


FIGURE 5.8. Absolute solvation free energy error for benzyl alcohol in kcal/mol as a function of homogenous scaling  $\alpha_\xi$  and field awareness parameter. The flux average and spread parameters are set to values of 2.0 and 0.5, respectively, so that only a couple of atoms are affected by the field-aware procedure.

aware parameters. In particular, as reported in Figure 5.8, simply by increasing the field-awareness of the model, one can achieve the same result as obtained by retuning of the SSCS model. The free energy of solvation of benzyl alcohol is overestimated by 19.66kcal/mol before reparameterization, reducing to an error of 10.11kcal/mol with  $\alpha_\xi = 0.98$ . The figure demonstrates how a similar solvation free energy can be achieved without decreasing  $\alpha_\xi$  and instead with a field-aware function tuned properly. In fact the figure seems to suggest that we could use a larger field awareness parameter. However, by visual inspection of the interface function for the best set of parameters (Figure 5.9, it appears that the improvement in solvation free energy for this compound is obtained by significantly reducing the size of






FIGURE 5.9. Benzyl Alcohol. Unscaled interface (left) represents the spherical functions as calculated by standard SSCS. Scaled interface (right) is based off an optimized set of parameters as reported in Table 5.1.

the soft-sphere around the oxygen atom. Increasing even further this effect would lead to an unphysically small solvation radius. The remaining discrepancy in solvation free energies should thus be ascribed to the aromatic portion of the molecule and to the assumptions of the model used to describe non-electrostatic interactions.

#### 5.5.4. Water

As in the previous systems, we consider the prototypical cases of the aqueous solvation of water and its derived ionic species, hydroxyl anion and hydronium cation. Although free energy of solvation of a solvent molecule in its own solution represents a pathological application of continuum solvation models, water is an interesting test case, which performs well in its neutral state, but remarkably poorly in its singly charged ionic state due to its small size. Even after the standard scaling that the SSCS performs, while the absolute solvation error from neutral water is only 0.81kcal/mol, the errors from the cation and anion are 18.56kcal/mol and 12.13kcal/mol respectively. The maximum flux cutoff is chosen to be close to the minimum flux of water, resulting in a general scaling of all atoms. In the different panels of Figures 5.10, we can see that a linear relation exists between the increase in the  $\alpha_\xi$  scaling factor and the corresponding optimal value of field awareness  $f_0$ . While

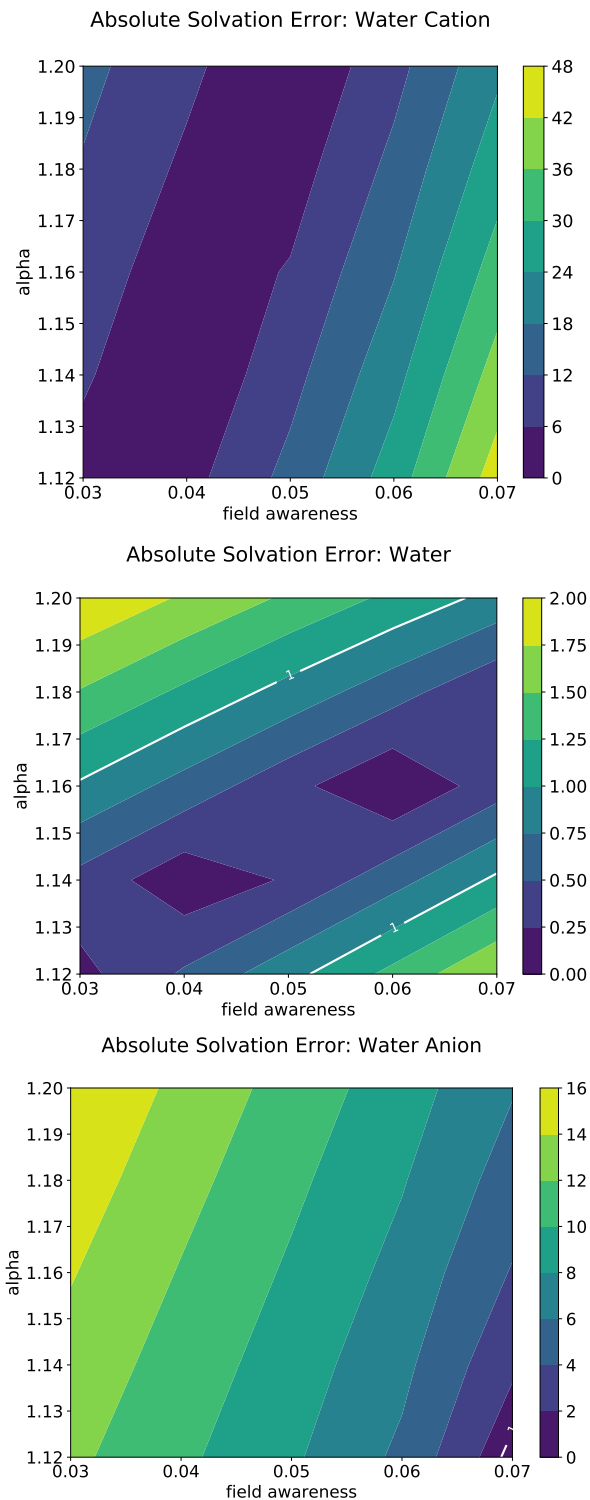


FIGURE 5.10. Absolute solvation free energy errors of neutral water (central panel), hydronium ion (top panel), and hydroxyl anion (bottom panel) in kcal/mol as functions of field awareness parameter and homogenous scaling parameter  $\alpha_\xi$ . The asymmetry parameter is set to zero, +1, and -1 for the three systems, respectively. Mean and spread values of the flux thresholds are set to 1.5 and 0.6.

the slope of this relationship is the same for cation and anion, the slope for neutral molecule is divided by factor of 4 due to the chosen values of the asymmetry parameter ( $\kappa = 0$  for neutral vs.  $\kappa = \pm 1$  for cation/anion) that introduce an additional scaling in the field-aware strength. Increasing the field-awareness from zero improves the solvation energy with respect to the experimental value, and as expected, the optimal field awareness parameter is different for all three cases. This disagreement in optimal field awareness can be overcome by an appropriate choice of an asymmetry parameter common to the three systems. One possible method for choosing the asymmetry parameter is by combining the optimal field awareness parameter individually obtained for cation ( $f_c$ ) and anion ( $f_a$ ) for an asymmetry of +1 and -1, respectively. The maximum shrinking effect obtained in the two systems is thus

$$(5.18) \quad \max_{\Phi} f_a t(\Phi) (-1 - \text{sgn}(\Phi))^2 = 4f_a,$$

$$(5.19) \quad \max_{\Phi} f_c t(\Phi) (1 - \text{sgn}(\Phi))^2 = 4f_c.$$

By imposing that the optimal global values of  $\kappa$  and  $f_0$  reproduce both effects, we obtain a system of two equations in two unknowns,

$$(5.20) \quad f_0(\kappa - 1)^2 = 4f_a,$$

$$(5.21) \quad f_0(\kappa + 1)^2 = 4f_c.$$

Given the above conditions, the determination of  $\kappa$  is trivial

$$(5.22) \quad \kappa = \frac{1}{R - 1} (R + 1 \pm 2\sqrt{R}),$$

where  $R = f_c/f_a$ , and consequently  $f_0$  can be found. For the hydroxyl anion and hydronium cation, parameters that produced reliable solvation free energy estimates for both would be  $\kappa = -0.16$  and  $f_0 = 0.24$ . Simulations exploiting a common set of parameters, including the two derived above, produce a systematic improvement in the accuracy of the calculations for all water species ( $\Phi_{\min} = 0.1$ ,  $\Phi_{\max} = 2.9$ ). Solvation free energy errors of 0.55kcal/mol, 2.67kcal/mol, and 3.26kcal/mol are obtained for neutral, cationic, and anionic water, respectively.

## 5.6. Conclusions

In conclusion, a new field-aware approach for continuum interfaces is proposed and analyzed. The field-aware interface has analytical, smoothly varying derivatives, allowing for its use in accurate geometry optimization or in molecular dynamics simulations. By exploiting reciprocal-space techniques to compute the electric field and derived properties, the overhead linked to this non-local definition is substantially reduced. The field-aware interface requires a number of new parameters, which can be tuned to reproduce experimental data on solvation free energies of neutral and charged compounds. An analysis of the main parameters involved in the model and of their effects on solvation free energies of a significant set of test molecules was also presented. Significant improvements over the standard approaches are demonstrated for a specifically tuned set of parameters. Nonetheless, an optimal parameterization aimed at improving the description of both neutral and charged compounds will require a more systematic search over the parameter space. In these regards, it will be important to focus the search of the individual parameters on specifically selected results, in order to isolate the physical effects introduced by this model. While experimental data on solvation free energies of neutral and charged molecules will be the main ingredient of the parameterization, first-principles molecular dynamics simulations of salts dissolution can represent a second unbiased source of determination of some of the parameters of the model. Eventually, parameterization on the physical properties of electrochemical interfaces, e.g. the differential capacitance of noble metal substrates, will also represent a potential strategy for the tuning of the presented approach.

## CHAPTER 6

### MACHINE LEARNING

To summarize the tools we have so far, there are a hierarchy of model families that tackle atomic processes of different scales. The truly exact solution for any system is limited by computational power, and an analytical solution for more than two bodies is infeasible. Up until recently, our computational power has forced us to be exceptionally clever in how we define models, in the sense that understanding of the theory and subsequent developments of empirical models are critical in the progress of models like molecular dynamics, which are bound by its reliance on well defined force-fields, in DFT, which is bound by well defined pseudo-potentials, and in continuum solvation, which is limited by the explicit solvent interactions that do affect regions of interest. The hierarchy itself is important in the sense that there are systems that we might be interested in that can only viably be approached by a single family of models, and it is important to define a problem with the intended tool in mind.

These families in the past did connect to each other somewhat, DFT can generate better ab-initio forcefields for use in molecular dynamics simulations. Continuum solvation can explicitly represent a layer of solvent in order to better capture the specific interactions with a solute. However, this solution only works for an isolated problem, in general, if we wanted to perform DFT to prepare forcefields, we would have to do this for every single molecular dynamics simulation. The forcefields are not particularly transferable, and therefore simulation time becomes a concerning factor.

The modern solution is machine learning. We know that conceptually there must be some kind of mapping between families of models. Consider the degrees of freedom in an n-body system. If the problem were completely general, then we're out of luck, finding a solution is not trivial in the slightest. However, chemistry is not completely general. There are patterns, and therefore, one can reduce the degrees of freedom in the system and still achieve generalizability. However, finding these mappings is non-trivial. The role that

machine learning provides us with a solution to this is due to the ability for algorithms to convert large quantities of data into these mappings.

The use of machine learning to improve or speed up simulation was mainly driven by what data had been accumulated at that point in time. Large quantities of DFT simulations had been run and ab-initio molecular dynamics heavily relied on DFT simulations for its potential energy surfaces (PES). If somehow we were able to take all the simulations already run previously and train a model to infer the PES from the typical inputs for an MD simulation, the efficiency of such simulations would increase dramatically. Likewise, if we were able to interpolate between our vast quantity of DFT data, we could in principle fill in all the gaps of simulations not yet run.

In 2006, Behler and Parrinello published a now hugely cited paper on solving for the PES with neural networks trained on DFT data. This is not the first paper to address the issue, but it is, at the time of publication, the most successful. They propose a neural network design and a representation for the input that reduces the initial problem, and showcase it on bulk silicon. The result is a method that is highly generalizable, and, once trained, performs orders of magnitude faster than DFT, expanding the potential of ab-initio MD methods to much larger systems, for longer simulation times.

These new generation models have increased accuracy over time (assuming high-quality data), because they rely on the quantity of data able to be generated by lower-level methods (methods with higher accuracy but lower speed). There are also limitations to these models. If the data does not exist, generating it can be an arduous task. If we apply a neural network to an input that lies outside of the training space, in other words, an extrapolation to structures very different from those trained, the validity of the model breaks down.

The general design of a model relies heavily on using physical knowledge to simplify the problem as much as possible before applying a machine learning algorithm. For example, a common input is the position of each atom. This is typically represented in Cartesian coordinates, but molecular properties do not depend on the absolute positions of the atoms. By translating or rotating a molecule, we change the coordinate input without changing

anything about the output, and therefore the first step would be to represent the positions of the atoms in a way that removes this invariance. Behler and Parrinello solve this issue by representing their atomic positions in terms of what they call symmetry functions, each of which is a function of all the atoms.

Another design complication is that typically the inputs are atomic properties but we wish to learn about molecular properties. Behler and Parrinello approach this by treating each atom as its own neural network that produces some energy value (and force), and then summing the energy values to reach the total energy of the system. Another approach assumes that the property of interest is additive, and so one can sum the atomic contributions at the beginning and then fit the weights of the decided model. This idea however requires training data to be scaled correctly with respect to the systems of interest in order to avoid any extrapolation.

One major limitation of continuum solvation models is the ability to capture the specific interactions between the solute and the solvent. For example, water has a tendency to form hydrogen bonds with solutes that are polarizable, thus reducing the total energy of the system. One would expect such a system to be more favorable than what a continuum solvation model might suggest. Another example would be the way that solvents structure themselves around solutes. Even if we simplified the scenario to zero interaction between the solute and the solvent except for the inability for the solvent to enter the space of the solute, we would have to consider the shape of the solvent, and how well it packs amongst members of its own species, and amongst all sorts of possible solvent molecules. Ignoring this packing ability, and we arrive at the theory of hard spheres. Many solvents do not pack like spheres, or they have different interaction radii depending on how they optimize for electrostatics. Water for example forms a hydrogen bonding network, and therefore will attempt to rearrange around a solute while preserving a loosely crystal-like structure (depending on the environment).

Remove all interactions between solvent and solute except for this hard sphere and we get a cavity. A cavity with the same radius as water will allow water to pack effectively

around it, but a cavity with radius 1.5 times that of water will not be expected to pack as well. If the theory of hard spheres gives us a mapping between the shape of the cavity and the energy to create the cavity, these sterically driven effects will manifest as a correction factor to that energy.

Simulating the structural tendencies of water around molecular shaped cavities with molecular dynamics gives us a more practical view of what is going on. Molecular dynamics treats water explicitly which accounts for specific interactions with the caveat of heavy performance losses over continuum solvation. We don't want to resort to this every time specific interactions need to be considered, and therefore a representative data set can be chosen, simulations run with molecular dynamics, and machine learning can be utilized for a more general model that is simple and fast.

Starting with cavitation energy, a single term in the calculation of solvation energy, is perhaps not the easiest task to prepare for, simply due to the fact that reliable experimental data on cavitation energy is limited or not collated, and the amount of simulation time to generate sufficient data is particularly large. On the other hand, understanding the cavitation energy as an individual term has its advantages

Machine learning can be loosely defined as the automation of the building of models by means of large quantities of data. There is a scale of dimensionality when defining models that describe systems. Low dimensionality models are typically simple to understand in terms of how the input affects the output. Often it is clear that a problem is low dimensionality, and therefore we can restrict this to begin with. For example, the extension on a spring can, for a certain range, be shown to be linearly dependent on the force applied on that spring. One can verify this claim by assuming a simple linear model to test the relationship between extension and force, and then generating data experimentally. Finally some suitable learning algorithm is selected. In this case, linear regression is a good choice, since there are not many parameters that seem to affect the output, and the relationship is seen to be linear. Linear regression in one dimension assumes a relationship of the form



$$(6.1) \quad y = cx + b$$

where  $y$  is the output,  $x$  is the input, and  $b$  and  $c$  are some constants.  $c$  is the constant of proportionality and  $b$  is the intercept, the value that  $y$  takes when  $x$  is zero. This is somewhat dependent on the scaling of the input and output. Often input and output is scaled so that machine learning algorithms perform more effectively, and for the correct scaling of parameters one may further simplify the model by setting the intercept to zero. In the general case, both constants need to be determined. Given some data of length  $n$  with known input  $x$  and output  $y$ , each pair of data points (enough to solve for 2 unknowns), or tuple  $(x_i, y_i, x_j, y_j)$  will typically produce different values for  $b$  and  $c$ , despite them following an approximately linear trend. We therefore determine a ‘best fit’ model, which compromises between data points and minimizes some arbitrarily defined cost function. A common cost function is ordinary least squares, which minimizes the square of the difference between predicted value  $y' = cx + b$  and the actual value  $y$ .

For more complicated systems, for example, finding the potential energy surface of an arbitrary atomic configuration, one must consider a high-dimensional input. We also may start to consider higher order terms, and mixed terms. The set of all possible terms becomes large as we consider higher order terms. Consider even some 3-dimensional input (or 3 input features), with up to quadratic terms. Already a general model would be

$$(6.2) \quad y = \sum_i^3 c_i x_i + \sum_i^3 \sum_j^3 c_{ij} x_i x_j$$

which has 12 terms and therefore 12 parameters to fit. Generally, to consider non-linear dependencies without knowing which non-linear terms are important becomes very expensive with respect to the number of parameters.

Simple regression models work up to a point, but eventually we want the model to automate more of the learning process. We want the ability to consider a very large feature space, with many parameters to fit, but also want to be able to automatically determine

which features are not affecting the output, or to distinguish between features that have the largest impact on the output versus higher order correction terms.

There are variants on linear regression models that can account that can resolve the over-fitting problems that arise when considering a large feature space.

## CHAPTER 7

### MACHINE LEARNING CAVITATION ENERGY

Solvation can be thought of as a combination of independent effects, one of which is cavitation, the act of creating a non-interacting cavity in a solvent. Cavitation is a particularly significant contributor towards hydrophobic solvation, where the interaction between solvent and solute is weak and the interactions between solvent molecules are comparatively strong. Of particular interest is the introduction of a bio-molecule to water, a solvent that possesses a network of hydrogen bonds, bonds that are disrupted by the creation of cavities. Such a system can be simulated using molecular dynamics, but this is an expensive process. One alternative is to represent the solvent as a collection of hard spheres. The energy required to form a cavity in this solvent can then be calculated by considering the statistical mechanics of a cavity sphere within a collection of solvent spheres. Only the repulsive part of the interaction of the cavity and the solvent is considered, and the resulting model can be expressed as a polynomial [76].

Scaled particle theory, as this statistical mechanics approach is commonly known, is able to make good predictions of the cavitation energy of small molecules, but does not capture the significance of the choice of cavity shape for protein sized molecules. In part this is due to the fact that in reality the solvent molecules do not have a spherical shape, however primarily, as is the case for water, they can possess strong electrostatic interactions with each other. As a result, errors are introduced by ignoring these effects that scale with the size of the cavity. The limitations of this simplified model have been documented numerous times in recent decades [44, 91, 48].

To capture some of these additional effects without incurring the costs of running explicit solvent simulations for every cavitation calculation, one can take inspiration from successes that use machine learning techniques. Generally, performance of model frameworks can be improved by direct use of machine learning on a particular calculation step, or accuracy can be increased by a combined use of machine learning and a lower level model that

would typically not be viable towards the system sizes of interest. Behler and Parrinello [12] showed that density functional theory (DFT) simulations could be used to train potential energy surfaces for use in molecular dynamics.

Bartok and Csanyi [8, 7] formalized a smooth overlap of atomic position (SOAP) representation that transformed the Cartesian representation of atomic positions into irrotational descriptors where the position of an atom was described in terms of its neighbourhood. They then exploited the power of Gaussian Processes [9], a promising non-parametric machine learning model by considering the similarity between pairs of structures using the aforementioned descriptors. This investigation starts by considering the atomic features that are a precursor the SOAP definition as input features, and then developing features more suited for cavitation energy, taking inspiration from SOAP and continuum models like PCM [97] and the soft-sphere continuum description [31].

Molecular dynamics has also been used to generate training data for the improvement of higher level models such as coarse grained water models, which allows for simulations on scales that are traditionally unattainable by molecular dynamics simulations. Similarly, molecular dynamics can provide training data useful for implicit solvation models, since they include approximations that molecular dynamics can account for. In general, molecular dynamics is expensive, however by machine learning learning a mapping between known properties and the desired output, only the initial investment into training the model is required. A small number of publications have been released at around the time of this investigation. They focus on using machine learning to ..

The goal of this investigation is to take the discrepancy in treatment between explicit and implicit solvent models, and focus solely on cavitation energy, with the intention of later generalizing to other energy contribution terms. Cavitation energy can be approximated by a very simple linear model,

$$(7.1) \quad E = \gamma S$$

where  $S$  is the surface area of the cavity. A better approximation can be derived from statistical mechanics. By treating the cavity and the solvent as hard spheres, Pierotti [76] derived a polynomial solution for the cavitation energy as a function of sphere radii. The function is primarily influenced by a quadratic term, and contributions from powers larger than 3 are negligible. In reality, cavities and solvents are not simply spherical in shape, and solvent molecules do not distribute themselves fully randomly due to their electrostatics. We therefore expect that molecular dynamics will show a linear dependence with respect to surface area with a higher order corrections.

Implicit solvent models are capable of reaching chemical accuracy for small organic systems, but their successes are more limited in applications that rely on specific interactions. In the case of complex structures that

### 7.1. Scaled Particle Theory

The statistical mechanics of rigid spheres is an interesting problem that saw most of its theoretical development in the late 50s / early 60s by Reiss, Helfand, Lebowitz, Frisch, Casberg, and Praestgaard [45, 56, 79, 80]. Of particular significance is the work by Pierotti, who formulated the Gibbs free energy of creating a cavity in a fluid of hard spheres.

Consider  $N$  spherically symmetric molecules with a diameter  $a_1$ , exerting attractive forces consistent with some fluid volume  $V$ . Exclude some spherical region (a cavity) of radius  $r$ . The idea is to consider the probability of this cavity,  $p_0(r, \rho)$ , where  $\rho = N/V$ , the number density of the fluid. This probability would be given by

$$(7.2) \quad p_0(r, \rho) = \exp\left(-\frac{W(r, \rho)}{kT}\right)$$

$W$  is described as the reversible work to produce a cavity of radius  $r$ .

The scaled particle theory attempts to determine the LHS (and therefore  $W$ ) based on statistical mechanics and geometric arguments. Consider the probability of finding a molecular center in a spherical shell radius  $r$ ,

$$(7.3) \quad 4\pi r^2 \rho \mathbf{G}(r, \rho) dr$$

where  $\rho \mathbf{G}$  is the conditional probability that a molecular center is location in that spherical shell.

The probability that no molecule exists inside the sphere multiplied by the probability that no molecule exists in a shell on this spherical radius is

$$(7.4) \quad \begin{aligned} p_0(r + dr) &= p_0(r) + \left( \frac{\partial p_0(r)}{\partial r} \right) dr \\ &= p_0(r) (1 - 4\pi r^2 \rho \mathbf{G}(r, \rho) dr) \end{aligned}$$

or

$$(7.5) \quad \frac{\partial}{\partial r} \ln p_0(r) = -4\pi r^2 \rho \mathbf{G}(r, \rho).$$

This can be related to  $W$ ,

$$(7.6) \quad \beta W(r, \rho) = 4\pi \rho \int_0^r r^2 \mathbf{G}(r, \rho) dr$$

$\mathbf{G}$  can be appropriately expanded like

$$(7.7) \quad \sum_i \mathbf{G}_i(\rho) \left( \frac{1}{r} \right)^i$$

To obtain an exact relationship, some assumptions have to be made. Take the case  $r \leq \sigma_1/2$  (which relates to a point solute), only a single hard sphere can have its center in the spherical region of radius  $r$  such that the spheres do not overlap. The probability that a cavity exists is

$$(7.8) \quad p_0 = 1 - \frac{4}{3}\pi r^3 \rho$$

This produces the interesting result

$$(7.9) \quad \beta W_0(r, \rho) = \ln \left( 1 - \frac{4}{3}\pi r^3 \rho \right)$$

which represents the reversible work of introducing a point solute into the fluid. For a solute sphere diameter  $\sigma_2$ , the cavity has a size of  $r = (\sigma_1 + \sigma_2)/2$ .

Much work has gone into pinpointing the coefficients for  $\mathbf{G}$ . The result is an expression for  $W$ ,

$$(7.10) \quad W(r, \rho) = K_0 + K_1 r + K_2 r^2 + K_3 r^3$$

In the original development of scaled particle theory, these  $K$  constants were obtained by Taylor expanding  $W$ , so

$$K_0 = W_0$$

$$K_1 = W'_0$$

$$K_2 = \frac{1}{2}W''_0$$

$$K_3 = \frac{1}{6}W'''_0$$

The result is,

$$(7.11) \quad \beta W(\mathbf{R}, \rho) = -\ln(1 - y) + \left( \frac{3y}{1 - y} \right) \mathbf{R} + \left[ \frac{3y}{1 - y} + \frac{9}{2} \left( \frac{y}{1 - y} \right)^2 \right] \mathbf{R}^2 + \frac{\beta y P}{\rho} \mathbf{R}^3$$

where  $y = \pi\rho\sigma_1^3/6$  is the reduced number density, and  $\mathbf{R} = \sigma_2/\sigma_1$ . To simplify the problem of cavitation, the solvent is kept constant, and the cavity is modified. Hence, just the change with respect to  $\sigma_2$  is of interest,

In the case of water, scaled particle theory suggests that the work done to create a cavity scales like a quadratic function. Pierotti reports for water,  $y = 0.371$ , which results in  $K_0 = 0.463$ ,  $K_1 = 1.769$ ,  $K_2 = 3.335$ . The linear term in eq. 7.11 is smaller than the quadratic term, and the cubic term is negligible in comparison for reasonable pressure values. This implies that a somewhat good approximation for the cavitation energy based off hard spheres is,

$$(7.12) \quad W = \gamma S$$

where  $S$  is the surface area of the cavity and  $\gamma$  is some constant to be determined. This approximation applies remarkably well for systems simulated with continuum solvation models, which all exploit this result for the cavitation energy. However more complicated molecules will surely exhibit behaviours in a solvent that cannot be captured by scaled particle theory. Scaled particle theory does not account for the complex shapes that molecules can manifest as.

If we replace solvent molecules with atoms rather than spheres, and consider a regular spherical cavity, we arrive at a slightly different result to scaled particle theory. Molecular dynamics can be used to set up a system of water molecules surrounding an atomic centered cavity with no interactions between the solvent and the cavity other than a repulsive potential that emulates a hard vacuum interface. This is somewhat unphysical and the cavitation energy term is not a property to be isolated experimentally.

In order to compute the free energy of cavitation computationally, a free energy method needs to be selected. Thermodynamic integration is a somewhat rudimentary strategy, but it is regarded as robust. Combining the idea of a hard sphere potential with thermodynamic integration is non trivial, since in principle a hard sphere potential should



be a step function of infinite potential. Thermodynamic integration, however, relies on the definition of a finite potential in order to integrate from 0 to  $V$ , where  $V$  is the potential of the step. Additionally, a step potential incurs instabilities in MD simulation and so resorting to a smooth potential is necessary.

Jover et al. [52] proposed a cut and shifted Mie potential with high powers in an effort to capture the hard potential in a way that

Water is described using TIP4P. In comparison to a hard sphere model of the solvent, water molecules now have electrostatic interactions with each other, they have a non-spherical shape, and they form an electrostatic dipole. The result is that water will structure itself in a way that is driven by the dipole dipole interactions, something that is not seen in the hard-sphere model.

The result is a drift away from the hard-sphere model as we move to larger cavities. The water will loosely structure itself around the spherical cavity in a way that favors certain cavity sizes over others.

## 7.2. Methods

### 7.2.1. General Framework

In order to generate training data for machine learning models, investigations often rely on a rich supply of pre-existing simulations. In the case of cavitation energy, it is a more challenging problem for two reasons. Firstly, a mechanism for calculating the free energy of cavitation for a system is not straightforward and requires expensive free energy molecular dynamics simulations, and secondly, it is not feasible to directly measure the cavitation energy experimentally. One must instead, indirectly infer the energy from subtracting known energy contributions from say, the solvation energy. As a consequence, we use the scaled particle theory developed by Pierotti in order to validate the results of the simulations. An exploratory approach was taken regarding the design of a machine learning based model. As with previous approaches, the position coordinates of the atoms would be used as training descriptors, but would have to be transformed into some coordinate independent representation first.

### 7.2.2. Descriptors

Descriptors for machine learning condensed matter simulations have a significant impact on the effectiveness of the parameterization of the final model. Since obtaining training data can be computationally expensive, it is preferential to minimize what the model needs to learn. For example, if our input features consist of the collection of atomic coordinates represented in Cartesian space, we need the model to infer that a translation of the molecule does not affect a property like solvation energy. Hence, one of the first efforts has been to recast the positions of atoms into a coordinate free representation system, one that is unchanged by both translations and rotations of the molecule of interest.

We consider the idea that cavitation is the sum of two effects, the volume displaced plus the ability for the water to structure around the cavity. The descriptors should be non-local, since the shape of the cavity can best be captured from the positions of the atoms by considering their position relative to one another. A coordinate free representation is typically non-local, since the reference points for the positions of atoms need to be relative to the molecule rather than the space in which it lies. We consider a popular representation for atomic positions in machine learning, SOAP descriptors [7] as a starting point towards learning cavitation energy. SOAP descriptors are non-local, and effectively describe the surrounding environment around each atom. This should give us a powerful distinction between different cavity shapes. However, it is expected to not scale particularly well for larger sized cavities, since the descriptors do not account for the fact that cavitation energy does not in principle, weight atomic contributions equally. The ability for water to structure around the cavity is determined solely by atoms exposed to the interface with the solvent molecules.

We derive a new family of descriptors inspired by the successes of SOAP and the inherent capabilities of continuum models to derive properties like surface area and atomic contributions to surface area, something that is not trivially available in SOAP. Consider the soft-sphere interface as described in eq 2.12. We can use this atomic based interface function along with functions for the quantum surface (eq. 2.17) and the quantum volume

(eq: 2.15) to define molecular descriptors. Now consider a spherical probe  $P(r)$  centered on an atom. It is possible to describe the contribution that an atom has on the surface area of the molecule by taking the same function used to compute the quantum surface but restricting the integration to the volume defined by the spherical probe. By varying the radius of this probe, it is possible to define a collection of descriptors that range from zero when the radius of the probe is smaller than the radius of the atom (note that in practice, since the functions are smooth, the quantity computed is non-zero), all the way to the total surface area of the molecule for a probe radius that fully engulfs the molecule. We label these as Environ surface descriptors. Similarly a collection of volume descriptors can be defined, labelled Environ volume descriptors.

$$(7.13) \quad S_P(r) = \int_{P(r)} \|\nabla s(\mathbf{r})\| d\mathbf{r}$$

$$(7.14) \quad V_P(r) = \int_{P(r)} s(\mathbf{r}) d\mathbf{r}$$

It is perhaps expected that a combination of these descriptors is important in closely representing the cavitation energy of a molecule. The volume is important in determining how much water needs displacing, whereas the surface descriptors should effectively represent the solvent-solvent interactions that are possible around the cavity. We generate a collection of Environ surface and volume descriptors with a varying size of probe from 0.1Å to 10.0Å. In order to filter these descriptors down to a set of useful descriptors, CUR decomposition is used [50].

### 7.2.3. Molecular Dynamics

A general system for measuring cavitation energy is comprised of a single solute surrounded by water molecules in a periodic cell, such that the solute does not interact with itself. To simulate water, the TIP4P model is chosen due to its simplistic implementation whilst still being able to accurately capture its interactions. The solute has no interactions




FIGURE 7.1. Visualization of Environ surface descriptor. Atomic centered spherical probe shown in green overlaps boundary defined by Environ soft-sphere interface (shown as a transparent 'bubble' surrounding the protein molecule). Overlap shown in bright green.

other than its ability to exert a repulsive force on atoms that cross a certain distance threshold. Ideally this force would be defined by an infinite step potential, so that an interaction is zero outside the cutoff and any incoming atoms will be reflected without any change of

kinetic energy. Defining such a potential in molecular dynamics is non-trivial since any function needs to be differentiable in order to compute for forces and maintain a stable simulation. Instead a steep but smooth potential is chosen. The Mie potential is the general form for a pairwise potential with attractive and repulsive terms. Each term is raised to an arbitrary power, set by two parameters that define the potential. A cut and shifted Mie potential is a piece wise potential that ensures a smooth potential that is only repulsive (the shift), and past some cutoff, exactly zero interaction (the cut).

In order to compute the free energy of cavitation from molecular dynamics simulations, thermodynamic integration was the method of choice [33],

$$(7.15) \quad \Delta F(O \rightarrow A) = \int_0^1 \langle U_A(\lambda) - U_O(\lambda) \rangle_\lambda d\lambda$$

where in context to cavitation energy, A is the state at which a cavity exists in solution and O is the state without a cavity in solution.  $\langle \dots \rangle_\lambda$  denotes the ensemble average for a system with potential energy function  $U(\lambda)$ .

Compared to other free energy methods, TI is robust, but the cost of calculation can scale quite poorly if the potential function is not well behaved. In the case of cavitation energy this issue does arise, however we will show that it is possible to interpolate via a variable power law instead of the standard linear or fixed power interpolation used in integration methods.

Applied to cavitation free energy, TI is not particularly well defined due to the fact that it requires taking the difference between two configurations, one of which is nonphysical. If we consider the zero potential case, molecules can freely move between the barrier. However once they are in such a configuration, the full potential must be considered, and this results in very high energy values that break computational models.

Another way to see this complication is in looking at the shape of the integrand function, which behaves somewhat similar to the function  $y = x^{-1}$ , for the cut and shifted Mie function that we wish to use for the cavity. This function is non-analytic when integrated

from zero, but we know that the free energy of cavitation is a finite value. One way to approach this is to consider a simple power law,  $U = \alpha\lambda^\gamma$  between two points, A and B. If

$$(7.16) \quad U_A = \alpha\lambda_A^\gamma$$

$$(7.17) \quad U_B = \alpha\lambda_B^\gamma$$

then

$$(7.18) \quad \alpha = \frac{U_B - U_A}{\lambda_B^\gamma - \lambda_A^\gamma}$$

$$(7.19) \quad \gamma = \frac{\log(U_B) - \log(U_A)}{\log(\lambda_B) - \log(\lambda_A)}$$

$$(7.20) \quad U_{A \rightarrow B}(\lambda) = \alpha\lambda^\gamma$$

$$(7.21) \quad \Delta F(A \rightarrow B) = \frac{\alpha}{\gamma + 1}(\lambda_B^{\gamma+1} - \lambda_A^{\gamma+1})$$

As  $\lambda \rightarrow 0$ , it is important that  $\gamma(\lambda)$  is less than one (which implies that the function does not diverge). Analysis from running MD simulations for  $\lambda = 10^{-1}$  all the way down to  $\lambda = 10^{-40}$  for varying powers of 10, shows that the power does indeed stay less than 1 and the function converges with each successive power of 10 contributing less and less to the total free energy beyond a certain  $\lambda$  threshold.

Consider a potential Mie( $n$ ,  $m$ ) where  $n$  is the power law governing the repulsive potential and  $m$  is the power law governing the attractive potential,

$$(7.22) \quad U_{n,m}(r) = \begin{cases} \epsilon U_{n,m}^0(r) & r < \sigma \left(\frac{n}{m}\right)^{\frac{1}{n-m}} \\ 0 & r \geq \sigma \left(\frac{n}{m}\right)^{\frac{1}{n-m}} \end{cases}$$

where

$$(7.23) \quad U_{n,m}^0(r) = \left( \frac{n}{n-m} \left(\frac{n}{m}\right)^{\frac{m}{n-m}} \left[ \left(\frac{\sigma}{r}\right)^n - \left(\frac{\sigma}{r}\right)^m \right] + 1 \right)$$

We wish to calculate  $\Delta F(O \rightarrow A)$  where  $O$  is a system with no potential and  $A$  is the system with some  $\text{Mie}(n, m)$  potential. Therefore,

$$(7.24) \quad \langle U_A(\lambda) - U_O(\lambda) \rangle_\lambda = \langle U_A(\lambda) \rangle_\lambda = \langle \lambda U_{n,m}(r) \rangle_\lambda$$

In order to calculate the free energy, we can take values of  $\lambda$  in the range  $(0, 1)$  and use a numerical integration scheme. Since the  $\text{Mie}(n, m)$  potential has a constant multiplicative factor built in  $(\epsilon)$ , we can repurpose that as the scaling factor for thermodynamic integration. We can print out the energy at each timestep, and then rescale it to get the energy that the system would have if the potential were not scaled by the  $\lambda$  parameter.

We finally arrive at the expression,

$$(7.25) \quad G_{\text{cav}} = \Delta F(O \rightarrow A) = \sum_{\lambda=\lambda_{\text{min}}}^{\lambda=1} \langle \lambda U_{n,m}(r) \rangle_\lambda$$

In order to construct a cavity, a set of atomic radii needs to be decided on. Since we treat the cavity as a purely repulsive force field without any electrostatic effects, it makes sense to use van der Waals radii [14].

#### 7.2.4. Machine Learning

To create a reference, we compare various linear regression models with the one dimensional linear model given by eq. 7.1.

### 7.3. Computational Details

Molecular Dynamics simulations are performed in DL POLY 4 [90] due to its reputation of being a well structured codebase that development oriented researchers can exploit to modify the standard operation of the code. This is often useful when investigating properties that aren't often looked at. To set up a free energy calculation, a number of parameters need to be optimized for. The number of water molecules needs to be such that adding extra molecules to the periodic cell (for constant density) does not significantly alter the configuration energy. By simulating varying sizes of cell, it is possible to find the point in which

the equilibrium energy of the system has converged with the number of water molecules simulated.

Since free energy requires statistical sampling, it is also necessary to find a compromise between convergence of a calculation and computational time taken to perform said calculation. There are two dimensions in play, the number of points to be integrated, and the amount of time a simulation is run for (which gives the number of sample points available). Since each simulation step has a causation link with its predecessor and successor step, for true random sampling, not all the points can be considered. Another way of looking at this effect is by considering taking the average of all the steps, and then taking this causation phenomenon into account in the error analysis by performing the method of block averages.

In order to prepare systems for production runs, equilibration is run and the complexity of this process mirrors the complexity of the system.

### 7.3.1. Small Molecules

For small molecules, a single equilibration simulation is performed for each lambda point, since these simulations are relatively cheap. The initial cells are prepared roughly by placing the water molecules on vertices in a three dimensional grid in Cartesian coordinates and removing those that are too close to cavity centers depending on the van der Waals radii of the atoms constituting each cavity. Equilibration is performed, first at constant pressure with a Berendsen barostat (1.0ps relaxation) and thermostat (0.5ps relaxation), with momentum re-sampling every 3fs, and temperature re-scaling every 5fs and with energy optimization after 10ps. The remainder of the equilibration is performed without strict momentum re-sampling and temperature re-scaling, starting from the optimized configuration found earlier for 0.19ns at 300K.

Production is run at constant volume with a Hoover thermostat (0.5ps relaxation) and a time-step of 1fs. Simulations are run for 0.2ns at each lambda point and the first 10ps are ignored due to instabilities from the switch of ensembles. Simulations are run at 300K with a van der Waals cutoff of 10 Angstrom. Water is modeled using TIP4P with electrostatics and the cavity is modeled by a cut and shifted Mie(30, 29) without electrostatics.



### 7.3.2. Proteins

For the protein cavities, water molecules were added by using the Packmol tool [62]. For each cavity, two initial simulations were performed sequentially. The first in constant energy, constant volume with a timestep of 0.5fs for 25ps. The second at constant volume, constant temperature (300K) with a timestep of 1fs for 0.5ns. The equilibration was performed per lambda point for 0.4ns at room temperature.

We then vary the two body potential (following thermodynamic integration) between the solute atoms (constituting the cavity) and the solvent atoms. The only information given for solute atoms are the position and the radii, the latter is chosen to match the van der Waals radius of each atom.

## 7.4. Results

### 7.4.1. Convergence Testing for Molecular Dynamics

The number of water atoms required to solvate a molecular dynamics system was chosen by considering increasing factors of 2 from 256 to 4096 and selecting a compromise between convergence for a cavitation energy calculation and time. The ability of DL POLY to run in parallel depends on the system size and cutoff distance. This means that the time for a simulation does not monotonically increase with respect to the number of water molecules in the system. Figure 7.2 suggests that for small water molecules solvated with more than 1500 water atoms, the cavitation energy is close to converged.

The first  $n$  steps of the simulation were expected to be unsuitable for analysis due to inherent instabilities in switching ensembles from equilibration to production. This can clearly be seen in figure 7.3, instabilities exist at the beginning of the simulation. The choice of where to apply the cutoff is made at 10ps.

In order to improve simulation accuracy, lambda points need to be chosen reasonably. Additionally, a compromise between simulation time, number of lambda points, and length of production run needs to be made. Error analysis is non-trivial to determine since exact cavitation energies are unknown, and with thermodynamic integration, the ideal case (ignoring

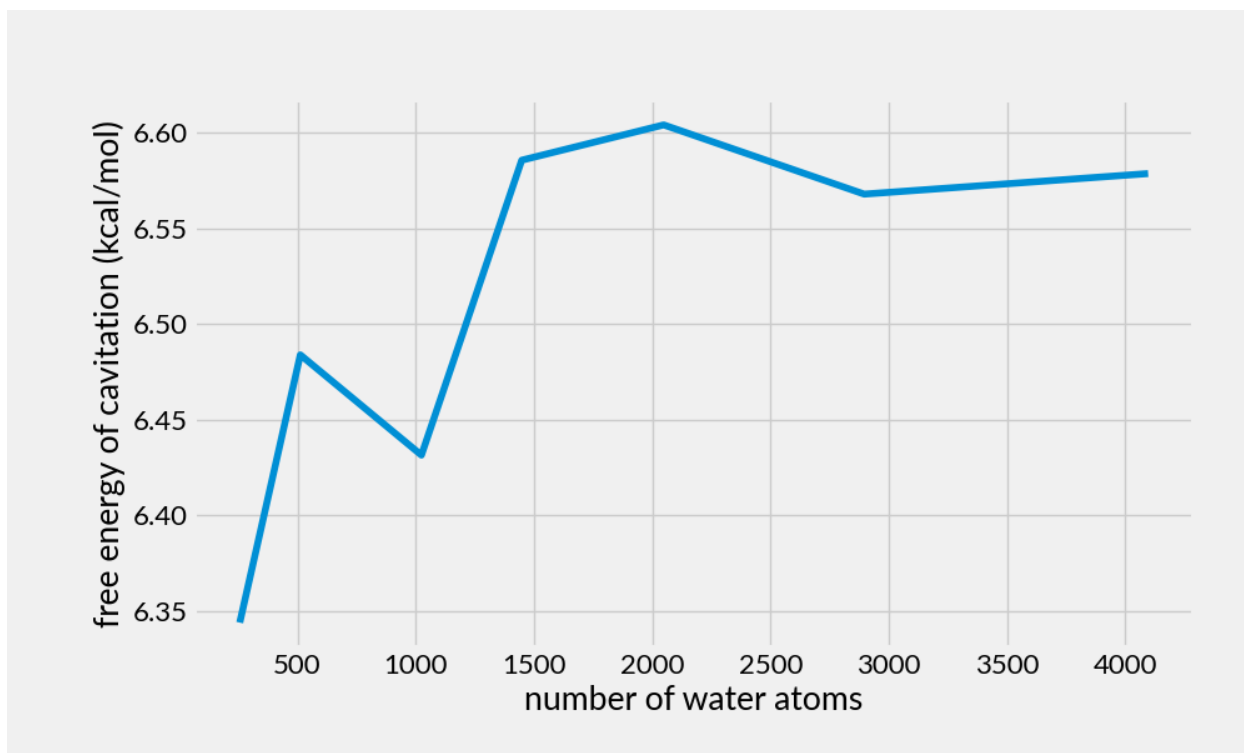


FIGURE 7.2. Free energy of cavitation of a water-like cavity immersed in a solution of water atoms. Data obtained as in section 7.3.1, with a production time of 1ns and a reduced van der Waals cutoff for simulations with smaller cells, or less water molecules.

computation time) would be running infinite points for an infinite length of time. Analysis on small molecular cavities shown in figure 7.4 suggests that increasing lambda points from 10 to 20 has a more significant impact on the convergence of the cavitation energy than the simulation time. Errors for the cavitation energy can be estimated by considering all of these convergence tests in combination. Taking the difference of free energy between 14 and 20 lambda points as within 0.05 kcal/mol, the difference of free energy between 1500 and 4000 water molecules as within 0.05 kcal/mol, and the convergence with respect to time as less than 0.05 kcal/mol after 200ps, and assuming these are independent sources suggests a total error of the order of 0.1 kcal/mol. The assumption here is that the convergence of cavitation energy or the gradient with respect to a convergence parameter will continue to decrease monotonically as the parameters continue to increase beyond what is tested for.

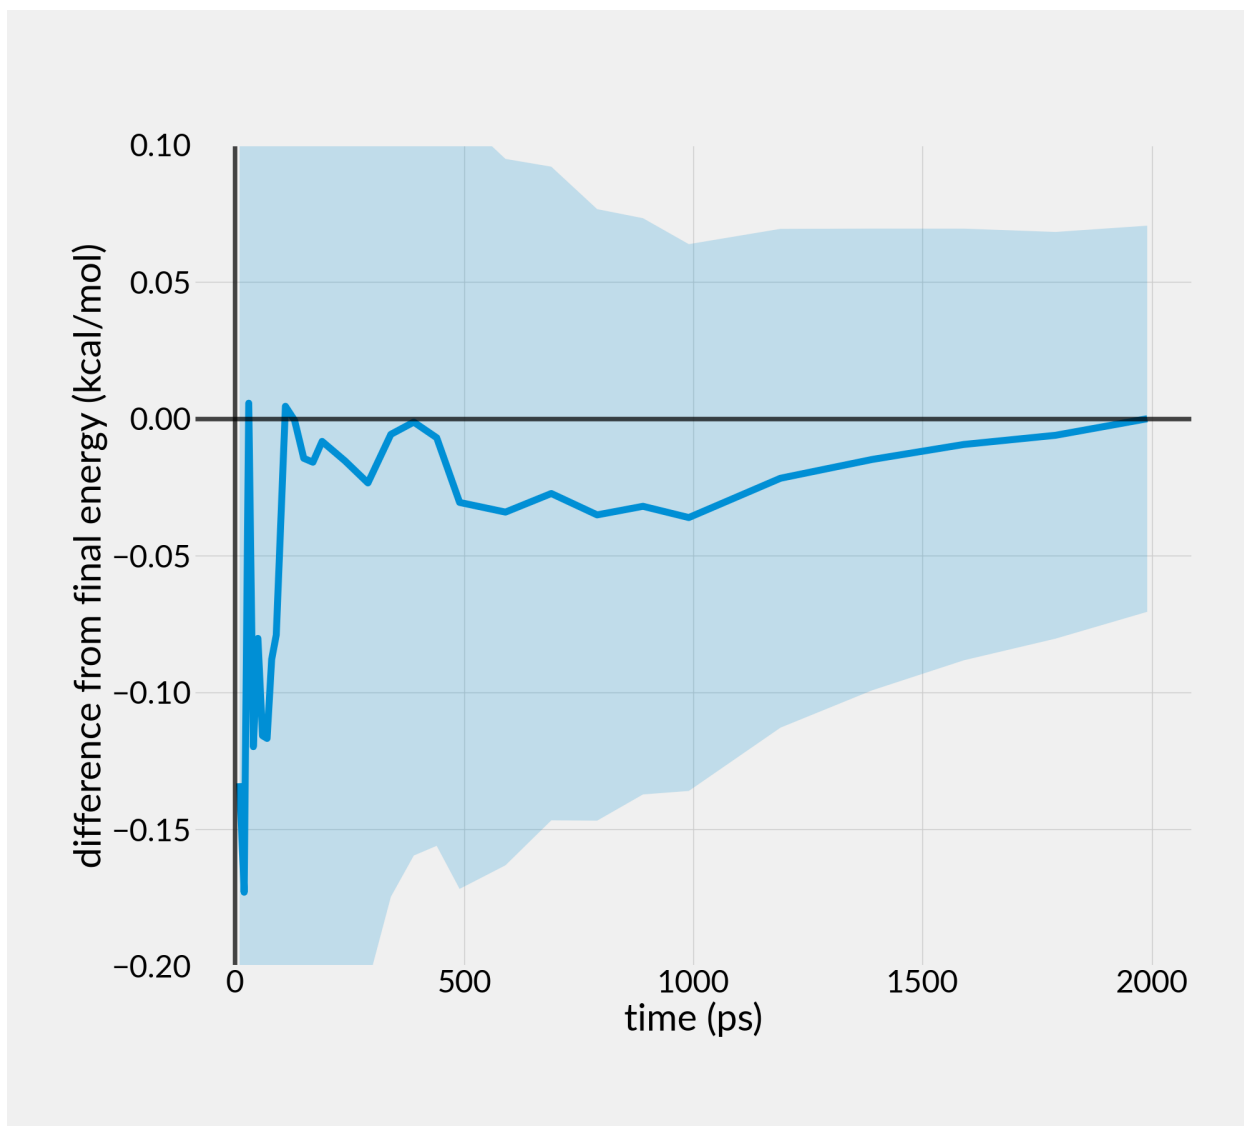


FIGURE 7.3. Free energy of cavitation reported as a difference with respect to the final energy. Plot represents the averaged cavitation energy of seven small molecular cavities taken with respect to the cavitation energy calculation at a snapshot after 2ns of simulation time. Shaded blue region shows a standard deviation of the free energy fluctuations that occur during simulations, aggregated as described in section B.

Another calculation for the error is accounted for in figure 7.3. Here, the shaded blue region represents one standard deviation of the free energy fluctuations that occur during simulations. For a discrete number of lambda points,

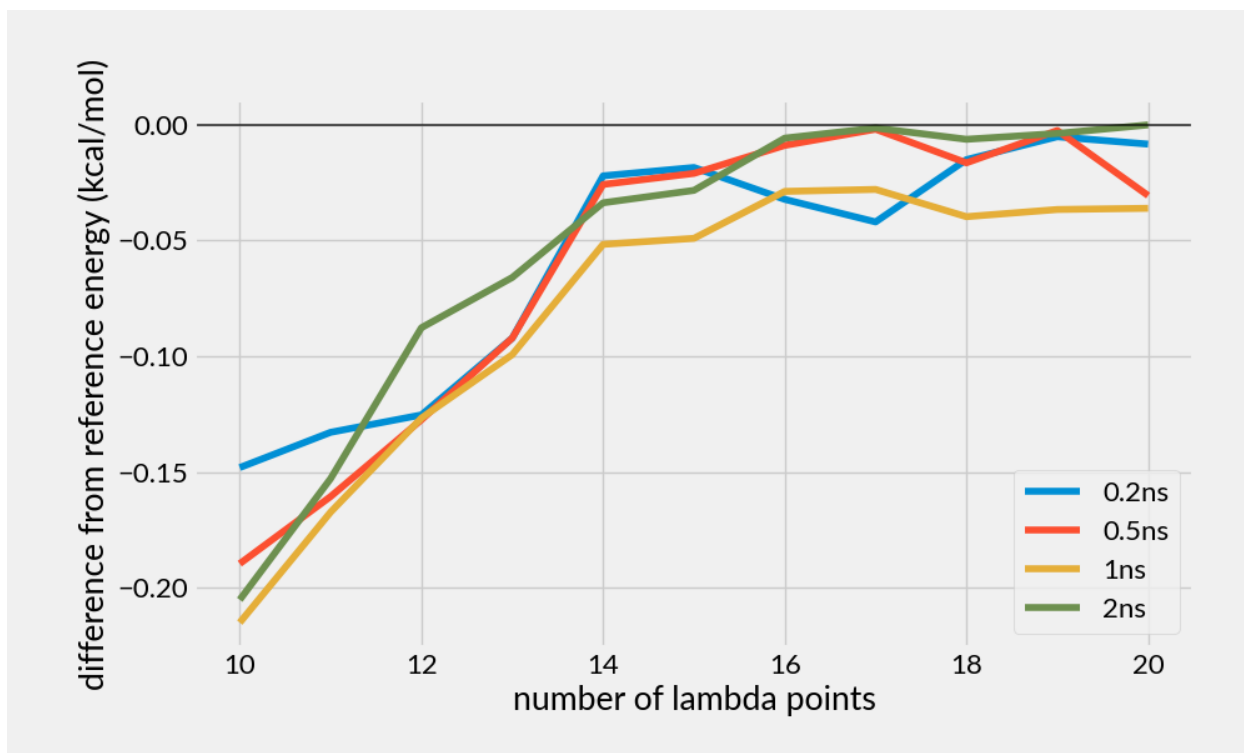


FIGURE 7.4. Free energy of cavitation reported as a different with respect to the final energy as calculated with 20 lambda points. Plots represent the averaged cavitation energy of seven small molecular cavities taken with respect to the cavitation energy calculation at a snapshot after 2ns of simulation time. This figure demonstrates how the clustering of the trendlines (convergence with respect to total time available for sampling) is smaller than the convergence over lambda points.

#### 7.4.2. Cavitation Energy of Spheres

In order to validate cavitation energy calculations, comparisons with scaled particle theory are made. By considering both the solvent molecules and the cavity has hard spheres one can derive equation 7.11 from statistical mechanics.

In the molecular dynamics simulations, the water molecules are represented by a four point model which have dipole interactions. We therefore expect to see a divergence between theory and the molecular dynamics calculations due to the specific interactions of the water molecules with each other as they arrange themselves around the cavity.




FIGURE 7.5. Work done to create a cavity diameter  $\sigma_2$  in TIP4P water in yellow, against hard-sphere equivalent 7.11 in blue ( $y = 0.371$ ). Pressure chosen to match simulations,  $T = 300K$ . Red line denotes quadratic term in scaled particle theory to demonstrate how the contributions come into play.

Figure 7.5 suggests that for atomic sized cavities, there will be relatively little difference between theory and simulation. It is somewhat tricky to compare a cavitation energy calculation with scaled particle theory. Scaled particle theory assumes rightly that the radius of water is approximately  $1.4\text{\AA}$ , this is due to the fact that water is a polarizable solvent and electrostatics reduces the mean distance between water molecules such that the effective radius is reduced. However when interacting with a hard cavity such that only cavitation effects are included and electrostatic effects are not, the radius of water to be used should be the van der Waals radius, which is approximately  $1.7\text{\AA}$ [58]. However there are still some takeaways to be made for such a comparison. Firstly, for cavities of equivalent size to the

solvent, the cavitation energy from molecular dynamics simulations agree closely with scaled particle theory. For small cavities, the shape fits a quadratic shape, which agrees with the idea that the quadratic term is dominating in the scaled particle theory expression [75]. There is a divergence for larger cavities. This is somewhat expected simply due to the fact that the steric effects of water should start to show as the cavity size increases. Additionally what can be seen is that certain sizes of cavities are more suited for the packing of water than others, since the gradient of the cavitation energy curve is seen to increase in a non-monotonic fashion.

Suppose we now take the molecular dynamics results for small organic molecules and fit them to scaled particle theory. The immediate issue is that these small molecular cavities are not spherical, and therefore an effective radius needs to be computed for each cavity. Treating cavities of this size is perhaps fine in the sense that the complex shapes that might significantly impact the structuring of water molecules around a cavity are limited by the small cavity sizes. Colominas [23] takes the approach of defining the radius of a molecule in two possible ways, either by rearranging the equation for the surface area of a sphere and deriving the radius from the surface area of the molecule (SPT-S), or by rearranging the equation for the volume of a sphere and deriving the radius from the volume of the molecule (SPT-V). Both of these properties can be calculated from the soft-sphere interface function. A comparison is then made with the Claverie-Pierotti formula for the cavitation energy around a cavity made up from a collection of hard spheres [21],

$$(7.26) \quad G_{\text{cav}} = \sum_{k=1}^N \frac{A_k}{4\pi R_k^2} G_{\text{cav}}(R_k)$$

which is a summation over all atoms in a molecule.  $A$  is the surface of the atom that is exposed. Essentially, this formula takes the cavitation energy for a hard sphere of radius equal to each atom and scales this based off its contribution to the surface area of the molecule. This should in principle scale well for larger molecules which have atoms not in contact with solvent molecules. There are two ways to calculate the  $G_{\text{cav}}(R)$ . Either by

	SPT-V	SPT-S	C-SPT	C-SPT-MD
$y$	0.335	0.3191	0.378	N/A
$\alpha$	1.0	1.0	1.02	1.123
MAE (SM) (kcal/mol)	0.481	0.395	0.836	1.024
MAE (P) (kcal/mol)	42.275	133.45	272.34	226.38

TABLE 7.1. The performance (given in terms of the mean absolute error with respect to molecular dynamics simulations) of scaled particle theory as optimized for small molecules (SM) is shown. These optimized parameters  $y$ ,  $\alpha$  are then applied to a set of proteins (P).

using Pierotti’s scaled particle theory to find the cavitation energy of a solvent of radius  $R$  (C-SPT), or by means of thermodynamic integration to find the cavitation energy of a cavity of radius  $R$  (C-SPT-MD). The Environ surface descriptors can be used to determine  $A$ , but since they are defined as error functions that smoothly switch at distance  $R$  from the atomic centres, it is necessary to scale these distances up by a small amount so as to fully capture the surface contribution per atom. This scaling factor is shown in table 7.1 as the parameter  $\alpha$ .

Out of the box, these models do not perform well when matched up against molecular dynamics simulations. This may well be due to the  $y$  parameter from scaled particle theory (influenced by the density of the solvent and the radius of the hard sphere representation of the solvent), for reasons explained earlier. Instead we decide to stretch these models as far as possible and optimize for the mean absolute error while freely varying the  $y$  parameter for each model. This parameter is shown in table 7.1 where applicable. We then compare these results with other models for the cavitation energy.

Table 7.1 shows that for small molecules, SPT-S performs the best, closely followed by SPT-V. These results agree with a similar analysis by Colominas [23]. C-SPT performs worse, and using MD results in place of the cavitation energy for single spheres does not improve things. When this best performing parameter set is applied to proteins, we interestingly see

that the SPT-V performs better than SPT-S. This can be somewhat rationalized by the fact that proteins are larger structures that have surfaces which may not be exposed to solvent. Surface irregularities will increase the effective radius of the protein cavity beyond what is representative.

### 7.4.3. Linear Regression on Descriptors

Since we know that there is a decent proportionality relationship between surface area and cavitation energy, it serves as a good validation tool and benchmark goal to perform analysis on surface descriptors. We also introduce usage of Environ-based surface descriptors here. These descriptors provide more localized information about the shape of the cavity, and therefore should be an improvement to the single surface area descriptor.

Ordinary Least Squares is used to fit the standard model,

$$(7.27) \quad \mathbf{y} = \mathbf{X}\beta + \epsilon$$

where  $\mathbf{y}$  represents the output descriptors, in this context, the MD simulated cavitation energy values.  $\mathbf{X}$  represents the input descriptors, in this context, the surface area terms (one feature per molecule).

Figure 7.6 shows the ability for quantum surface alone to describe cavitation energy. For the referenced continuum solvation models [3, 31], the cavitation term assumes no constant correction factor (the  $\epsilon$  in equation 7.27). We initially consider only small molecules, performing ordinary least squares regression with 5-fold cross validation. If  $\epsilon = 0$ , then the MAE is calculated to be 0.646. If  $\epsilon \neq 0$ , then the MAE is calculated to be 0.3297. Such a simple expression for the cavitation energy does indeed work effectively for small molecules, however, applying this result directly to proteins incurs errors that are much more significant (126 kcal/mol assuming  $\epsilon \neq 0$ ). Clearly, relying on surface area alone as a predictor is insufficient as we scale up.

Adding the proteins to the training set results in a MAE of 17.567 kcal/mol for predicting proteins and we have sacrificed our ability to predict small molecules as effectively.



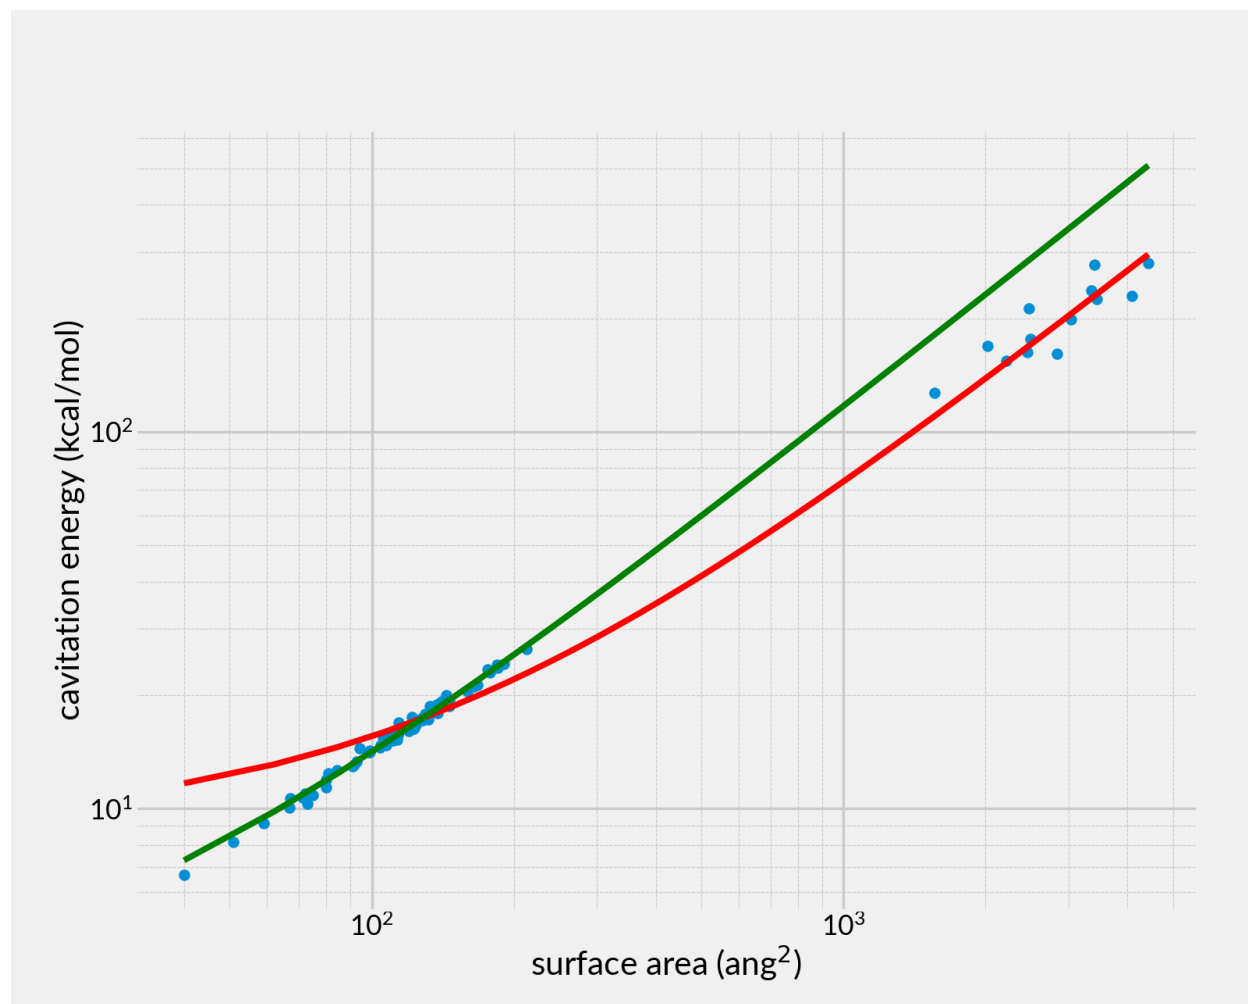


FIGURE 7.6. The linear relationship between cavitation energy and surface area. Surface area generated by quantum surface calculation with soft-sphere interface. Red trend line generated by ordinary least squares regression with a combined training set of small molecules and proteins with 3-fold cross validation. Green trend line generated by ordinary least squares regression with the small molecules used for training data.

The ideal model would predict small molecules with high accuracy, and still perform well for proteins. It is also apparent that while the spread for small molecules is small enough for a linear model with respect to the surface area to capture cavitation effects, the higher spread for proteins suggests that there are factors at play that do not derive from surface area.

The quantum volume is not far behind in its ability to describe the cavitation energy

in simple terms for small molecules. Here,  $\epsilon$  is found to be near zero, and the MAE is calculated to be 0.4768 kcal/mol. The quantum volume is a worse predictor of protein cavitation energy (184kcal/mol) than the quantum surface. Combining these two is more powerful than considering them separately, which is to be expected since physically they provide different information about the cavity. MAE for small molecules is found to be 0.3138 kcal/mol and for proteins, the MAE is 103 kcal/mol. As seen in the results from the quantum surface, adding proteins to the training data improves performance for proteins but reduces performance for the small molecules.

The table 7.2 summarizes the results from ordinary least squares regression. For the training data available, the results suggest that if the goal is a single model for the cavitation energy of both small molecules and proteins, trying to minimize the error with proteins as part of the training set does impact the performance on small molecules quite significantly. Arguably the best performing set of descriptors appears to be the quantum volume plus Environ surface descriptors. If we train on small molecules, the performance is better than what the quantum surface descriptor alone is capable of, and the MAE (15.84 kcal/mol) when extrapolating this model towards proteins is an improvement over all other descriptor choices (assuming that no proteins are used for training).

The table also showcases some of the drawbacks of ordinary least squares regression when applied to datasets with a large number of descriptors. Rather than relying on a better choice of descriptors, one can choose regression techniques with the inbuilt ability to avoid multicollinearity issues. Ridge, Lasso, and Elastic Net regression all remove possible correlations in descriptors. In an effort to improve performance further, all three techniques were used on a combination of descriptors. We report here the best performing combinations.

Improving on the accuracy of the small molecular cavitation energies is possible with both ridge and lasso regression, however the improvement is minimal and the resulting model performs worse when attempting to extrapolate towards proteins. Minimizing the error for proteins beyond what is attainable by simple regression is possible (down to around 10 kcal/mol), albeit at the cost of again reducing the performance for simple molecules.

	SM $\rightarrow$ SM	SM $\rightarrow$ P	SM + P $\rightarrow$ SM	SM + P $\rightarrow$ P
Quantum Surface (QS)	0.3297	125.7	1.261	17.57
Quantum Volume (QV)	0.4768	184.0	1.542	20.68
QS + QV	0.3138	102.7	0.7694	14.25
Environ Surface (ES)	0.7288	32.35	1.471	22.14
Environ Volume (EV)	0.8067	181.6	0.9615	23.47
ES + QS	0.2831	39.90	0.8988	21.55
ES + QV	0.2990	15.84	1.146	13.26
ES + QS + QV	0.2743	72.34	0.8715	14.59
EV + QS	0.2819	275.6	0.9096	12.98
EV + QV	0.3056	400.2	1.104	15.61
EV + QS + QV	0.2846	143.9	0.9228	13.31
ES + EV + QS	0.2792	482.6	0.7528	19.52
ES + EV + QV	0.2966	245.6	0.7987	19.31
ES + EV + QS + QV	0.2834	573.3	0.7361	21.94

TABLE 7.2. All values above are MAE (kcal/mol). A subset of descriptors was used for the Environ surface and Environ volume descriptors, obtained using CUR decomposition. Best performing results are obtained for each combination of descriptors for the small molecule set and then applied to proteins. Best performing results are obtained from a mix of small molecules and proteins towards proteins and then applied to small molecules.

This level of accuracy is quite encouraging considering the lack of training data for proteins obtained for this investigation, and the fact that the cavitation energy for proteins is typically of the order of  $10^2 - 10^3$ .

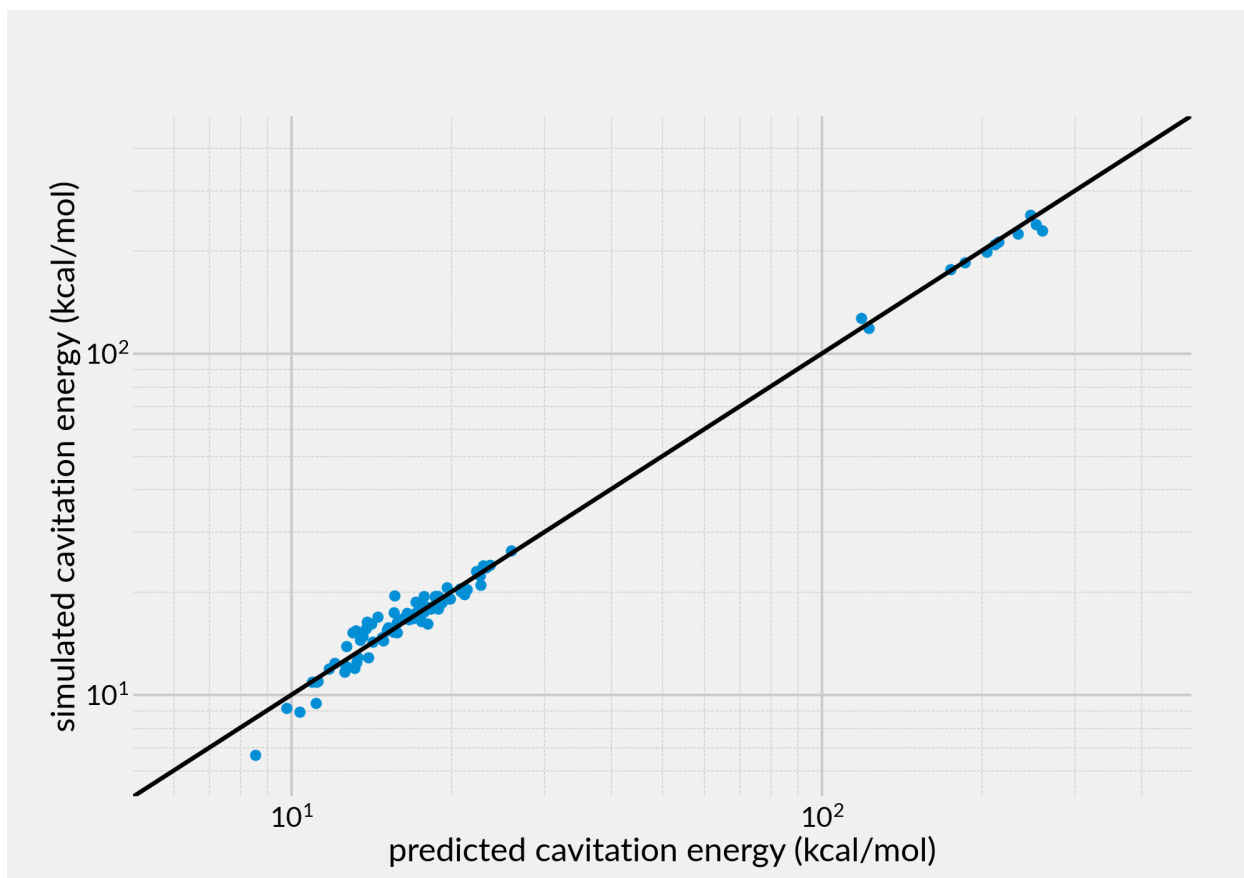


FIGURE 7.7. Best performing set of descriptors, Environ surface + volume, with lasso regression. Predicted cavitation energy plotted against simulated cavitation energy on a log-log graph. 3-fold cross validation used, graphic shows the final randomized testing set. MAE for small molecules 0.8 kcal/mol, MAE for proteins 10.2 kcal/mol.

APPENDIX A

FIELD AWARE DERIVATIVES

One of the key component of the analytic derivatives of the field-aware interfaces is represented by the derivatives of the field fluxes with respect to the electronic density and the ionic positions. The latter derivative is particularly cumbersome, as the dependence on ionic positions enters both the calculation of the field vector and the direction normal to the soft-sphere interface, which is expressed in terms of the soft-sphere gradient. Moreover, both terms depend explicitly on all the ionic positions, so that all cross terms need to be included, i.e. the partial derivative of the field flux through the soft sphere surrounding atom  $a$  with respect to the change in position of a second atom  $c$  is different from 0.

The derivative that we ought to express is the following

$$\begin{aligned}
\nabla_{\mathbf{R}_c} \Phi_a &= \nabla_{\mathbf{R}_c} \int \mathbf{E}(\mathbf{r}) \cdot \nabla h_a(r_a) \prod_{b \neq a} h_b(r_b) d\mathbf{r} \\
\text{(A.1)} \quad &= \int \nabla_{\mathbf{R}_c} [\mathbf{E}(\mathbf{r}) \cdot \nabla h_a(r_a)] \prod_{b \neq a} h_b(r_b) d\mathbf{r} \\
&+ \sum_b (1 - \delta_{ab}) \int [\mathbf{E}(\mathbf{r}) \cdot \nabla h_a(r_a)] \nabla_{\mathbf{R}_c} h_b(r_b) \prod_{d \neq a, b} h_d(r_d) d\mathbf{r}
\end{aligned}$$

where the factor  $(1 - \delta_{ab})$  entering the second term is introduced to enforce the fact that the index  $b$  must be different from  $a$ . To simplify the above expression, we can exploit the fact that for central functions we have

$$\text{(A.2)} \quad \nabla_{\mathbf{R}} f(\mathbf{r} - \mathbf{R}) = -\nabla f(\mathbf{r} - \mathbf{R})$$

or equivalently in this case,

$$\text{(A.3)} \quad \nabla_{\mathbf{R}} h(r) = -\nabla h(r).$$

The summation in the second term can be simplified by noting that the only term different from zero occurs when  $c$  and  $b$  are the same, namely

$$\begin{aligned}
\text{(A.4)} \quad \sum_b (1 - \delta_{ab}) \int [\mathbf{E}(\mathbf{r}) \cdot \nabla h_a(r_a)] \nabla_{\mathbf{R}_c} h_b(r_b) \prod_{d \neq a, b} h_d(r_d) d\mathbf{r} = \\
(1 - \delta_{ac}) \int [\mathbf{E}(\mathbf{r}) \cdot \nabla h_a(r_a)] \nabla_{\mathbf{R}_c} h_c(r_c) \prod_{d \neq a, c} h_d(r_d) d\mathbf{r} =
\end{aligned}$$

$$(\delta_{ac} - 1) \int [\mathbf{E}(\mathbf{r}) \cdot \nabla h_a(r_a)] \nabla h_c(r_c) \prod_{d \neq a, c} h_d(r_d) d\mathbf{r}$$

The first integral in Eq. (A.1), instead, can be simplified by using the vector calculus identity for the differential operator applied to a vector scalar product, namely

$$\begin{aligned} \nabla(\mathbf{A} \cdot \mathbf{B}) &= \nabla(A_x B_x + A_y B_y + A_z B_z) \\ (A.5) \quad &= (\nabla \otimes \mathbf{A})\mathbf{B} + (\nabla \otimes \mathbf{B})\mathbf{A} \\ &= \mathbf{A} \times (\nabla \times \mathbf{B}) + \mathbf{B} \times (\nabla \times \mathbf{A}) + (\mathbf{A} \cdot \nabla)\mathbf{B} + (\mathbf{B} \cdot \nabla)\mathbf{A} \end{aligned}$$

where the notation  $\mathbf{u} \otimes \mathbf{v}$  is adopted to indicate the outer product between two vectors. By using the above identity, the quantity  $\nabla_{\mathbf{R}_c} [\mathbf{E}(\mathbf{r}) \cdot \nabla h_a(r_a)]$  can thus be recast into the following expression

$$\begin{aligned} (A.6) \quad \nabla_{\mathbf{R}_c} [\mathbf{E}(\mathbf{r}) \cdot \nabla h_a(r_a)] &= [\nabla_{\mathbf{R}_c} \otimes \mathbf{E}(\mathbf{r})] \nabla h_a(r_a) + [\nabla_{\mathbf{R}_c} \otimes \nabla h_a(r_a)] \mathbf{E}(\mathbf{r}) \\ &= -[\nabla \otimes \nabla \phi(z_c \delta(r_c))] \nabla h_a(r_a) - \delta_{ac} [\nabla \otimes \nabla h_a(r_a)] \mathbf{E}(\mathbf{r}) \\ &= -\mathbf{H}[\phi(z_c \delta(r_c))] \nabla h_a(r_a) - \delta_{ac} \mathbf{H}[h_a(r_a)] \mathbf{E}(\mathbf{r}) \end{aligned}$$

In the second passage of the above derivation we expressed the electric field originated by the ion  $c$  in terms of the gradient of its electrostatic potential

$$(A.7) \quad \mathbf{E}_c(\mathbf{r}) = \nabla \phi(z_c \delta(r_c))$$

while, in the last passage we recast the outer products in terms of the Hessian matrix

$$(A.8) \quad \nabla \otimes \nabla f = \begin{pmatrix} \frac{\partial^2 f}{\partial x^2} & \frac{\partial^2 f}{\partial x \partial y} & \frac{\partial^2 f}{\partial x \partial z} \\ \frac{\partial^2 f}{\partial y \partial x} & \frac{\partial^2 f}{\partial y^2} & \frac{\partial^2 f}{\partial y \partial z} \\ \frac{\partial^2 f}{\partial z \partial x} & \frac{\partial^2 f}{\partial z \partial y} & \frac{\partial^2 f}{\partial z^2} \end{pmatrix} = \mathbf{H}[f].$$

By combining the above relations, the partial derivative of the field flux with respect to the atomic positions can be written as

$$(A.9) \quad \nabla_{\mathbf{R}_c} \Phi_a = - \int (\mathbf{H}[\phi(z_c \delta(r_c))] \nabla h_a(r_a) + \delta_{ac} \mathbf{H}[h_a(r_a)] \mathbf{E}(\mathbf{r})) \prod_{b \neq a, c} h_b(r_b) d\mathbf{r}$$

$$+ (\delta_{ac} - 1) \int [\mathbf{E}(\mathbf{r}) \cdot h_a(r_a)] \nabla h_c(r_c) \prod_{b \neq a, c} h_b(r_b) d\mathbf{r}$$



APPENDIX B  
ERROR ANALYSIS

## Block Averages

Due to the computational effort of obtaining the time correlation function, we usually estimate statistical errors in static quantities by studying the behaviour of block averages, defined over a finite time  $t_B$ ,

$$(B.1) \quad \bar{A}_B \equiv \frac{1}{t_B} \int_0^{t_B} dt A(t)$$

In the context of a simulation with discrete time, we calculate a block mean for each finite time block and then take the variance of this value. As the block time increases, this variance is expected to converge to a constant value indicating that there is no longer any correlation between blocks.

In practice, we start from a block size of 1, calculating and storing the variance, and then continuing the increase the block size, until the last  $n_v$  variances themselves have a standard deviation of less than 10% (an arbitrary stop point) of the most recent variance value. The oldest variance value is then taken to be the variance of the distribution (being the first value in a series of variance values which have shown to be converged sufficiently). This estimate is quite liberal since it is important to keep these simulations from running for too long.

## Error Propagation

So far we have a collection of points that together can describe the cavitation energy of a system due to thermodynamic integration. The error of each of these values is calculated by the method of block averages, however some manipulation is required in order to get the error of the computed cavitation energy.

There are a few complications in the thermodynamic approach. This is explained more verbosely above, but in summary, the energy and the value of lambda follow an inverse power law. After some analysis, it seems as if this inverse power law is less than 1 and therefore the integration is expected to converge. In order to numerically integrate I decided to assume a simple power law for interpolating the function between points, since usually the

assumption is that the function between points follows close to some kind of polynomial or preferably, a straight line. Due to the asymptotic nature of the integration, such a method would vastly overestimate the actual energy.

Suppose we have a series of values  $\langle U(\lambda) \rangle$  for values of  $\lambda$  in the range  $(0, 1)$ , spaced in powers of 10. We calculate the power law between adjacent points,

$$(B.2) \quad U_i = a_i \lambda_i^{b_i}$$

$$(B.3) \quad U_{i+1} = a_i \lambda_{i+1}^{b_i}$$

$$(B.4) \quad E_i = \frac{a_i}{b_i + 1} (\lambda_{i+1}^{b_i+1} - \lambda_i^{b_i+1})$$

From the above equations,  $U_i$  is obtained from the molecular dynamics simulations,  $\lambda$  is set in advance and therefore it is possible to calculate  $a$  and  $b$ . Finally

$$(B.5) \quad E = \sum_{i=0}^{n-1} E_i$$

From the method of block averages, we can calculate a value of the variance of each  $U_i$ ,  $\sigma_{U_i}$  and then for the variance of the cavitation energy, we use standard error propagation techniques (assuming that the points are independent). Since

$$(B.6) \quad b_i = \frac{\log \frac{U_{i+1}}{U_i}}{\log \frac{\lambda_{i+1}}{\lambda_i}}$$

$$(B.7) \quad \sigma_{b_i} = \frac{\sqrt{\left(\frac{\sigma_{U_{i+1}}}{U_{i+1}}\right)^2 + \left(\frac{\sigma_{U_i}}{U_i}\right)^2}}{\left| \log \left( \frac{\lambda_{i+1}}{\lambda_i} \right) \right|}$$

$$(B.8) \quad a_i = U_i \lambda_i^{-b_i}$$

$$(B.9) \quad \sigma_{a_i} = U_i \lambda_i^{-b_i} \sqrt{(\sigma_{b_i} \log \lambda_i)^2 + \left(\frac{\sigma_{U_i}}{U_i}\right)^2}$$

## REFERENCES

- [1] O Andreussi, N Hörmann, F Nattino, and N Marzari, *Solvent-Aware Interfaces in Continuum Solvation*, J. Chem. Theory Comput. 15 (2018), no. 3, 1996–2009.
- [2] O Andreussi, F Nattino, I Dabo, I Timrov, G Fisicaro, Goedecker S., and N Marzari, *Environ*, [www.quantum-environment.org](http://www.quantum-environment.org), 2018.
- [3] Oliviero Andreussi, Ismaila Dabo, and Nicola Marzari, *Revised self-consistent continuum solvation in electronic-structure calculations*, J. Chem. Phys. 136 (2012), no. 6.
- [4] Oliviero Andreussi, Nicholas Georg Hormann, Francesco Nattino, Giuseppe Fisicaro, Stefan Goedecker, and Nicola Marzari, *Solvent-aware Interfaces in Continuum Solvation*, Journal of Chemical Theory and Computation (2019), [acs.jctc.8b01174](https://doi.org/10.1021/acs.jctc.8b01174).
- [5] Oliviero Andreussi and Nicola Marzari, *Electrostatics of solvated systems in periodic boundary conditions*, Phys. Rev. B 90 (2014), no. 24, 245101.
- [6] Oliviero Andreussi, Francesco Nattino, Ismaila Dabo, Iurii Timrov, Giuseppe Fisicaro, Stefan Goedecker, and Nicola Marzari, *Environ 1.0: a library for environment effect in first-principles simulations of materials*, Environ.1.0 (2018), [www.quantum-environment.org](http://www.quantum-environment.org).
- [7] Albert P. Bartók, James Kermode, Noam Bernstein, and Gábor Csányi, *Machine learning a general-purpose interatomic potential for silicon*, Phys. Rev. X 8 (2018), 041048.
- [8] Albert P. Bartók, Risi Kondor, and Gábor Csányi, *On representing chemical environments*, Phys. Rev. B 87 (2013), 184115.
- [9] Albert P. Bartók and Gábor Csányi, *Gaussian approximation potentials: A brief tutorial introduction*, International Journal of Quantum Chemistry 115 (2015), no. 16, 1051–1057.
- [10] Artem Baskin and David Prendergast, *Improving Continuum Models to Define Practical Limits for Molecular Models of Electrified Interfaces*, J. Electrochem. Soc. 164 (2017), no. 11, 3438–3447.

- [11] Martin Z Bazant, Mustafa S Kilic, Brian D Storey, and Armand Ajdari, *Towards an understanding of induced-charge electrokinetics at large applied voltages in concentrated solutions*, Adv. Colloid Interface Sci. 152 (2009), no. 1-2, 48–88.
- [12] Jörg Behler and Michele Parrinello, *Generalized Neural-Network Representation of High-Dimensional Potential-Energy Surfaces*, Phys. Rev. Lett. 98 (2007), no. 14, 146401.
- [13] P Maarten Biesheuvel, Alejandro A Franco, and Martin Z Bazant, *Diffuse Charge Effects in Fuel Cell Membranes*, J. Electrochem. Soc. 156 (2009), no. 2, B225.
- [14] A. Bondi, *van der waals volumes and radii*, The Journal of Physical Chemistry 68 (1964), no. 3, 441–451.
- [15] Nicéphore Bonnet and Nicola Marzari, *First-principles prediction of the equilibrium shape of nanoparticles under realistic electrochemical conditions*, Phys. Rev. Lett. 110 (2013), no. 8, 86104.
- [16] Itamar Borukhov, David Andelman, and Henri Orland, *Steric effects in electrolytes: A modified poisson-boltzmann equation*, Phys. Rev. Lett. 79 (1997), 435–438.
- [17] ———, *Adsorption of large ions from an electrolyte solution: a modified Poisson-Boltzmann equation*, Electrochim. Acta 46 (2000), no. 2-3, 221–229.
- [18] Quinn Campbell and Ismaila Dabo, *Electrochemical stability and light-harvesting ability of silicon photoelectrodes in aqueous environments*, The Journal of Chemical Physics 151 (2019), no. 4, 044109.
- [19] Quinn Campbell, Daniel Fisher, and Ismaila Dabo, *Voltage-dependent reconstruction of layered  $\text{bi}_2\text{wo}_6$  and  $\text{bi}_2\text{moo}_6$  photocatalysts and its influence on charge separation for water splitting*, Phys. Rev. Materials 3 (2019), 015404.
- [20] David Leonard Chapman, *LI. A contribution to the theory of electrocapillarity*, Philos. Mag. 25 (1913), no. 148, 475–481.
- [21] P. Claverie, *Intermolecular interactions: from diatomics to biopolymers*, B. Pullman, Wiley, New York (1978), 69–305.
- [22] Matteo Cococcioni, Francesco Mauri, Gerbrand Ceder, and Nicola Marzari, *Electronic-*

- Enthalpy Functional for Finite Systems Under Pressure*, Phys. Rev. Lett. 94 (2005), no. 14, 145501.
- [23] Carles Colominas, F.Javier Luque, Jordi Teixidó, and Modesto Orozco, *Cavitation contribution to the free energy of solvation.: Comparison of different formalisms in the context of mst calculations*, Chemical Physics 240 (1999), no. 1, 253–264.
- [24] I Dabo, E Cancès, Y L Li, and N Marzari, *Towards First-principles Electrochemistry*, arXiv (2008), 0901.0096.
- [25] Ismaila Dabo, *Towards first-principles electrochemistry*, Ph.D. thesis, Massachusetts Institute of Technology, 2008.
- [26] Ismaila Dabo, Yanli Li, Nicéphore Bonnet, and Nicola Marzari, *Ab Initio Electrochemical Properties of Electrode Surfaces*, Fuel Cell Science (A Wicowski and J Norskov, eds.), John Wiley & Sons, 2010, pp. 415–431.
- [27] C Dupont, O Andreussi, and N Marzari, *Self-consistent continuum solvation (SCCS): the case of charged systems.*, J. Chem. Phys. 139 (2013), no. 21, 214110.
- [28] Jean-Luc Fattebert and François Gygi, *First-principles molecular dynamics simulations in a continuum solvent*, Int. J. Quantum Chem. 93 (2003), no. 2, 139–147.
- [29] Jean-Luc L Fattebert and François Gygi, *Density functional theory for efficient ab initio molecular dynamics simulations in solution.*, Journal of Computational Chemistry 23 (2002), no. 6, 662–666.
- [30] G Fisticaro, L Genovese, O Andreussi, N Marzari, and S Goedecker, *A generalized Poisson and Poisson-Boltzmann solver for electrostatic environments*, J. Chem. Phys. 144 (2016), no. 1, 14103.
- [31] Giuseppe Fisticaro, Luigi Genovese, Oliviero Andreussi, Sagarmoy Mandal, Nisanth N Nair, Nicola Marzari, and Stefan Goedecker, *Soft-Sphere Continuum Solvation in Electronic-Structure Calculations*, J. Chem. Theory Comput. 13 (2017), no. 8, 3829–3845.
- [32] Alessandro Fortunelli, William A Goddard, Yao Sha, Ted H Yu, Luca Sementa, Giovanni Barcaro, and Oliviero Andreussi, *Dramatic increase in the oxygen reduction re-*

- action for platinum cathodes from tuning the solvent dielectric constant*, *Angew. Chem. Int. Ed Engl.* 53 (2014), no. 26, 6669–6672 (en).
- [33] Daan Frenkel and Berend Smit, *Understanding molecular simulation: From algorithms to applications*, second ed., Computational Science Series, vol. 1, Academic Press, San Diego, 2002.
- [34] C L Fu and K M Ho, *External-charge-induced surface reconstruction on Ag(110)*, *Phys. Rev. Lett.* 63 (1989), no. 15, 1617–1620.
- [35] Kevin F Garrity, Joseph W Bennett, Karin M Rabe, and David Vanderbilt, *Pseudopotentials for high-throughput DFT calculations*, *Comput. Mat. Sci.* 81 (2014), 446–452.
- [36] Kevin L Gering, *Prediction of electrolyte viscosity for aqueous and non-aqueous systems: Results from a molecular model based on ion solvation and a chemical physics framework*, *Electrochim. Acta* 51 (2006), no. 15, 3125–3138.
- [37] P Giannozzi, O Andreussi, T Brumme, O Bunau, M Buongiorno Nardelli, M Calandra, R Car, C Cavazzoni, D Ceresoli, M Cococcioni, N Colonna, I Carnimeo, A Dal Corso, S de Gironcoli, P Delugas, R A DiStasio, A Ferretti, A Floris, G Fratesi, G Fugallo, R Gebauer, U Gerstmann, F Giustino, T Gorni, J Jia, M Kawamura, H.-Y. Ko, A Kokalj, E Küçükbenli, M Lazzeri, M Marsili, N Marzari, F Mauri, N L Nguyen, H.-V. Nguyen, A Otero-de-la Roza, L Paulatto, S Poncé, D Rocca, R Sabatini, B Santra, M Schlipf, A P Seitsonen, A Smogunov, I Timrov, T Thonhauser, P Umari, N Vast, X Wu, and S Baroni, *Advanced capabilities for materials modelling with Quantum ESPRESSO*, *J. Phys.: Condens. Matter* 29 (2017), no. 46, 465901.
- [38] Paolo Giannozzi, Stefano Baroni, Nicola Bonini, Matteo Calandra, Roberto Car, Carlo Cavazzoni, Davide Ceresoli, Guido L Chiarotti, Matteo Cococcioni, Ismaila Dabo, Andrea Dal Corso, Stefano De Gironcoli, Stefano Fabris, Guido Fratesi, Ralph Gebauer, Uwe Gerstmann, Christos Gougoussis, Anton Kokalj, Michele Lazzeri, Layla Martinamos, Nicola Marzari, Francesco Mauri, Riccardo Mazzarello, Stefano Paolini, Alfredo Pasquarello, Lorenzo Paulatto, Carlo Sbraccia, Alexander Smogunov, Paolo Umari, Andrea Dal Corso, Stefano de Gironcoli, Stefano Fabris, Guido Fratesi, Ralph Gebauer,



- Uwe Gerstmann, Christos Gougoussis, Anton Kokalj, Michele Lazzeri, Layla Martin-  
 samos, Nicola Marzari, Francesco Mauri, Riccardo Mazzarello, Stefano Paolini, Alfredo  
 Pasquarello, Lorenzo Paulatto, Carlo Sbraccia, Sandro Scandolo, Gabriele Scлаuzero,  
 Ari P Seitsonen, Alexander Smogunov, Paolo Umari, and Renata M Wentzcovitch,  
*QUANTUM ESPRESSO: a modular and open-source software project for quantum  
 simulations of materials*, J. Phys.: Condens. Matter 21 (2009), no. 39, 395502.
- [39] Paolo Giannozzi, Stefano Baroni, Nicola Bonini, Matteo Calandra, Roberto Car, Carlo  
 Cavazzoni, Davide Ceresoli, Guido L Chiarotti, Matteo Cococcioni, Ismaila Dabo, An-  
 drea Dal Corso, Stefano De Gironcoli, Stefano Fabris, Guido Fratesi, Ralph Gebauer,  
 Uwe Gerstmann, Christos Gougoussis, Anton Kokalj, Michele Lazzeri, Layla Martin-  
 samos, Nicola Marzari, Francesco Mauri, Riccardo Mazzarello, Stefano Paolini, Al-  
 fredo Pasquarello, Lorenzo Paulatto, Carlo Sbraccia, Alexander Smogunov, and Paolo  
 Umari, *QUANTUM ESPRESSO: a modular and open-source software project for quan-  
 tum simulations of materials*, J. Phys.: Condens. Matter 21 (2009), 395502.
- [40] Paolo Giannozzi, Stefano Baroni, Nicola Bonini, Matteo Calandra, Roberto Car, Carlo  
 Cavazzoni, Davide Ceresoli, Guido L Chiarotti, Matteo Cococcioni, Ismaila Dabo,  
 Andrea Dal Corso, Stefano de Gironcoli, Stefano Fabris, Guido Fratesi, Ralph Gebauer,  
 Uwe Gerstmann, Christos Gougoussis, Anton Kokalj, Michele Lazzeri, Layla Martin-  
 Samos, Nicola Marzari, Francesco Mauri, Riccardo Mazzarello, Stefano Paolini, Alfredo  
 Pasquarello, Lorenzo Paulatto, Carlo Sbraccia, Sandro Scandolo, Gabriele Scлаuzero,  
 Ari P Seitsonen, Alexander Smogunov, Paolo Umari, and Renata M Wentzcovitch,  
*QUANTUM ESPRESSO: a modular and open-source software project for quantum  
 simulations of materials*, J. Phys.: Condens. Matter 21 (2009), no. 39, 395502.
- [41] M Gouy, *Sur la constitution de la charge électrique à la surface d'un électrolyte*, J.  
 Phys. Theor. Appl. 9 (1910), no. 1, 457–468.
- [42] Deniz Gunceler, Kendra Letchworth-Weaver, Ravishankar Sundararaman, Kathleen A  
 Schwarz, and T A Arias, *The importance of nonlinear fluid response in joint density-*

- functional theory studies of battery systems*, *Model. Simul. Mater. Sci. Eng.* 21 (2013), no. 7, 74005.
- [43] J. Hansen, R. Ruedy, M. Sato, and K. Lo, *GLOBAL SURFACE TEMPERATURE CHANGE*, *Reviews of Geophysics* 48 (2010), no. 4, RG4004.
- [44] Robert C Harris, Alexander H Boschitsch, and Marcia O Fenley, *Sensitivities to parameterization in the size-modified Poisson-Boltzmann equation*, *J. Chem. Phys.* 140 (2014), no. 7, 75102.
- [45] E. Helfand, H. Reiss, H. L. Frisch, and J. L. Lebowitz, *Scaled particle theory of fluids*, *The Journal of Chemical Physics* 33 (1960), no. 5, 1379–1385.
- [46] H Helmholtz, *Ueber einige Gesetze der Vertheilung elektrischer Ströme in körperlichen Leitern mit Anwendung auf die thierisch-elektrischen Versuche*, *Ann. Phys.* 165 (1853), no. 6, 211–233.
- [47] Nicolas G Hörmann, Oliviero Andreussi, and Nicola Marzari, *Grand canonical simulations of electrochemical interfaces in implicit solvation models*, *The Journal of Chemical Physics* 150 (2019), no. 4, 41730.
- [48] David M. Huang, Phillip L. Geissler, and David Chandler, *Scaling of hydrophobic solvation free energies*, *The Journal of Physical Chemistry B* 105 (2001), no. 28, 6704–6709.
- [49] Sebastiaan P. Huber, Spyros Zoupanos, Martin Uhrin, Leopold Talirz, Leonid Kahle, Rico Häuselmann, Dominik Gresch, Tiziano Müller, Aliaksandr V. Yakutovich, Casper W. Andersen, Francisco F. Ramirez, Carl S. Adorf, Fernando Gargiulo, Snehal Kumbhar, Elsa Passaro, Conrad Johnston, Andrius Merkys, Andrea Cepellotti, Nicolas Mounet, Nicola Marzari, Boris Kozinsky, and Giovanni Pizzi, *Aiida 1.0, a scalable computational infrastructure for automated reproducible workflows and data provenance*, *Scientific Data* 7 (2020), no. 1, 300.
- [50] Giulio Imbalzano, Andrea Anelli, Daniele Giofré, Sinja Klees, Jörg Behler, and Michele Ceriotti, *Automatic selection of atomic fingerprints and reference configurations for*

- machine-learning potentials*, The Journal of Chemical Physics 148 (2018), no. 24, 241730.
- [51] Ryosuke Jinnouchi and Alfred B Anderson, *Electronic structure calculations of liquid-solid interfaces: Combination of density functional theory and modified Poisson-Boltzmann theory*, Phys. Rev. B 77 (2008), no. 24, 245417.
- [52] J. Jover, A. J. Haslam, A. Galindo, G. Jackson, and E. A. Müller, *Pseudo hard-sphere potential for use in continuous molecular-dynamics simulation of spherical and chain molecules*, Journal of Chemical Physics 137 (2012), no. 14, 144505.
- [53] Naiwrit Karmodak and Oliviero Andreussi, *Catalytic activity and stability of two-dimensional materials for the hydrogen evolution reaction*, ACS Energy Letters 5 (2020), no. 3, 885–891.
- [54] Naiwrit Karmodak, Luca Bursi, and Oliviero Andreussi, *Oxygen evolution and reduction on two-dimensional transition metal dichalcogenides*, J. Phys. Chem. Lett. 13 (2022), no. 1, 58–65 (en).
- [55] James R Kermode, *f90wrap: an automated tool for constructing deep python interfaces to modern fortran codes*, Journal of Physics: Condensed Matter 32 (2020), no. 30, 305901.
- [56] J. L. Lebowitz, E. Helfand, and E. Praestgaard, *Scaled particle theory of fluid mixtures*, The Journal of Chemical Physics 43 (1965), no. 3, 774–779.
- [57] Kendra Letchworth-Weaver and T A Arias, *Joint density functional theory of the electrode-electrolyte interface: Application to fixed electrode potentials, interfacial capacitances, and potentials of zero charge*, Phys. Rev. B 86 (2012), no. 7, 75140.
- [58] A J Li and R Nussinov, *A set of van der waals and coulombic radii of protein atoms for molecular and solvent-accessible surface calculation, packing evaluation, and docking*, Proteins 32 (1998), no. 1, 111–127 (en).
- [59] Alexander Y Lozovoi and Ali Alavi, *Reconstruction of charged surfaces: General trends and a case study of Pt(110) and Au(110)*, Phys. Rev. B 68 (2003), no. 24, 245416.
- [60] R Mancinelli, A Botti, F Bruni, M A Ricci, and A K Soper, *Hydration of Sodium*,

- Potassium, and Chloride Ions in Solution and the Concept of Structure Maker/Breaker*, J. Phys. Chem. B 111 (2007), 13570–13577.
- [61] Shaun A. Marcott, Jeremy D. Shakun, Peter U. Clark, and Alan C. Mix, *A reconstruction of regional and global temperature for the past 11,300 years*, Science 339 (2013), no. 6124, 1198–1201.
- [62] L Martínez, R Andrade, E G Birgin, and J M Martínez, *PACKMOL: a package for building initial configurations for molecular dynamics simulations*, J. Comput. Chem. 30 (2009), no. 13, 2157–2164 (en).
- [63] Nicola Marzari, David Vanderbilt, Alessandro De Vita, and M C Payne, *Thermal Contraction and Disordering of the Al(110) Surface*, Phys. Rev. Lett. 82 (1999), no. 16, 3296–3299.
- [64] Marko Melander, Mikael Kuisma, Thorbjørn Christensen, and Karoliina Honkala, *Grand-Canonical Approach to Density Functional Theory of Electrocatalytic Systems: Thermodynamics of Solid-Liquid Interfaces at Constant Ion and Electrode Potentials*, ChemRxiv (2018).
- [65] Benedetta Mennucci, *Polarizable Continuum Model*, WIREs Computational Molecular Science 2 (2012), 386–404.
- [66] S. Miertuš, E. Scrocco, and J. Tomasi, *Electrostatic interaction of a solute with a continuum. a direct utilization of ab initio molecular potentials for the prevision of solvent effects*, Chemical Physics 55 (1981), no. 1, 117–129.
- [67] Matthew M. Montemore, Oliviero Andreussi, and J. Will Medlin, *Hydrocarbon adsorption in an aqueous environment: A computational study of alkyls on cu(111)*, The Journal of Chemical Physics 145 (2016), no. 7, 074702.
- [68] Nicolas Mounet, Marco Gibertini, Philippe Schwaller, Davide Campi, Andrius Merkys, Antimo Marrazzo, Thibault Sohler, Ivano Eligio Castelli, Andrea Cepellotti, Giovanni Pizzi, and Nicola Marzari, *Two-dimensional materials from high-throughput computational exfoliation of experimentally known compounds*, Nature Nanotechnology 13 (2018), no. 3, 246–252.

- [69] Yasuya Nakayama and David Andelman, *Differential capacitance of the electric double layer: The interplay between ion finite size and dielectric decrement*, J. Chem. Phys. 142 (2015), no. 4, 44706.
- [70] E R Nightingale, *Phenomenological Theory of Ion Solvation. Effective Radii of Hydrated Ions*, J. Phys. Chem. 63 (1959), no. 9, 1381–1387.
- [71] M Otani and O Sugino, *First-principles calculations of charged surfaces and interfaces: A plane-wave nonrepeated slab approach*, Phys. Rev. B 73 (2006), no. 11, 115407.
- [72] John P Perdew, Kieron Burke, and Matthias Ernzerhof, *Generalized Gradient Approximation Made Simple [Phys. Rev. Lett. 77, 3865 (1996)]*, Phys. Rev. Lett. 78 (1997), no. 7, 1396.
- [73] Pearu Peterson, *F2py: a tool for connecting fortran and python programs*, Int. J. Comput. Sci. Eng. 4 (2009), 296–305.
- [74] J Chem Phys, Francesco Nattino, Matthew Truscott, Nicola Marzari, Oliviero Andreussi, Matthew Truscott, and Nicola Marzari, *Continuum models of the electrochemical diffuse layer in electronic-structure calculations Continuum models of the electrochemical diffuse layer in electronic-structure calculations*, 041722 (2019), no. December 2018.
- [75] Robert A. Pierotti, *THE SOLUBILITY OF GASES IN LIQUIDS <sup>1</sup>*, The Journal of Physical Chemistry 67 (1963), no. 9, 1840–1845.
- [76] ———, *A scaled particle theory of aqueous and nonaqueous solutions*, Chemical Reviews 76 (1976), no. 6, 717–726.
- [77] G Prandini, A Marrazzo, I E Castelli, N Mounet, and N Marzari, *A Standard Solid State Pseudopotentials (SSSP) library optimized for accuracy and efficiency (Version 1.0, data download)*, 2018.
- [78] A K Rappe, C J Casewit, K S Colwell, W A Goddard, and W M Skiff, *Uff, a Full Periodic-Table Force-Field for Molecular Mechanics and Molecular-Dynamics Simulations*, J. Am. Chem. Soc. 114 (1992), no. 25, 10024–10035.
- [79] H. Reiss and R. V. Casberg, *Radial distribution function for hard spheres from scaled*

- particle theory, and an improved equation of state*, The Journal of Chemical Physics 61 (1974), no. 3, 1107–1114.
- [80] H. Reiss, H. L. Frisch, and J. L. Lebowitz, *Statistical mechanics of rigid spheres*, The Journal of Chemical Physics 31 (1959), no. 2, 369–380.
- [81] Stefan Ringe, Harald Oberhofer, Christoph Hille, Sebastian Matera, and Karsten Reuter, *Function-Space-Based Solution Scheme for the Size-Modified Poisson–Boltzmann Equation in Full-Potential DFT*, J. Chem. Theory Comput. 12 (2016), no. 8, 4052–4066.
- [82] Stefan Ringe, Harald Oberhofer, and Karsten Reuter, *Transferable ionic parameters for first-principles Poisson-Boltzmann solvation calculations: Neutral solutes in aqueous monovalent salt solutions*, J. Chem. Phys. 146 (2017), no. 13, 134103.
- [83] Hannah Ritchie and Max Roser, *Carbon dioxide and greenhouse gas emissions*, Our World in Data (2020), <https://ourworldindata.org/co2-and-other-greenhouse-gas-emissions>.
- [84] ———, *Energy*, Our World in Data (2020), <https://ourworldindata.org/energy>.
- [85] Helen K Roobottom, H Donald B Jenkins, Jack Passmore, and Leslie Glasser, *Thermochemical Radii of Complex Ions*, J. Chem. Educ. 76 (1999), no. 11, 1570.
- [86] D A Scherlis, J L Fattebert, F Gygi, M Cococcioni, and N Marzari, *A unified electrostatic and cavitation model for first-principles molecular dynamics in solution*, J. Chem. Phys. 124 (2006), no. 7, 74103.
- [87] Luca Sementa, Oliviero Andreussi, William A. Goddard III, and Alessandro Fortunelli, *Catalytic activity of pt38 in the oxygen reduction reaction from first-principles simulations*, Catal. Sci. Technol. 6 (2016), 6901–6909.
- [88] X Shao, O Andreussi, D Ceresoli, M Truscott, A Baczewski, Q Campbell, and M Pavanello, *QEpy*, <https://gitlab.com/shaoxc/qepy>, 2021.
- [89] Patrice Simon and Yury Gogotsi, *Materials for electrochemical capacitors*, Nat. Mater. 7 (2008), no. 11, 845–854.

- [90] W. Smith and I. T. Todorov, *A short description of dl\_poly*, *Molecular Simulation* 32 (2006), no. 12-13, 935–943.
- [91] Gabriele Cesare Sosso, Sebastiano Caravati, Grant Rotskoff, Suriyanarayan Vaikuntanathan, and Ali Hassanali, *On the role of nonspherical cavities in short length-scale density fluctuations in water*, *J. Phys. Chem. A* 121 (2017), no. 1, 370–380 (en).
- [92] T. Stein, *Carbon dioxide peaks near 420 parts per million at mauna loa observatory*, (2021).
- [93] Otto Stern, *Zur Theorie der Elektrolytischen Doppelschicht*, *Z. Elektrochem.* 30 (1924), 508.
- [94] Ravishankar Sundararaman, Kendra Letchworth-Weaver, and Kathleen A Schwarz, *Improving accuracy of electrochemical capacitance and solvation energetics in first-principles calculations*, *J. Chem. Phys.* 148 (2018), no. 14, 144105.
- [95] Ravishankar Sundararaman and Kathleen Schwarz, *Evaluating continuum solvation models for the electrode-electrolyte interface: Challenges and strategies for improvement*, *J. Chem. Phys.* 146 (2017), no. 8, 84111.
- [96] M N Tamashiro and H Schiessel, *Where the linearized Poisson-Boltzmann cell model fails: The planar case as a prototype study*, *Phys. Rev. E* 68 (2003), no. 6, 66106.
- [97] Jacopo Tomasi, Benedetta Mennucci, and Roberto Cammi, *Quantum mechanical continuum solvation models*, *Chemical Reviews* 105 (2005), no. 8, 2999–3093.
- [98] Sergio Trasatti, *Work function, electronegativity, and electrochemical behaviour of metals: Iii. electrolytic hydrogen evolution in acid solutions*, *Journal of Electroanalytical Chemistry and Interfacial Electrochemistry* 39 (1972), no. 1, 163–184.
- [99] M Truscott, A Jayanth, N Martinez, and O Andreussi, *AiiDA Environ*, <https://github.com/environ-developers/aaida-environ>, 2022.
- [100] Matthew Truscott and Oliviero Andreussi, *Field-aware interfaces in continuum solvation*, *The Journal of Physical Chemistry B* 123 (2019), no. 16, 3513–3524, PMID: 30943719.
- [101] Zachary W. Ulissi, Aayush R. Singh, Charlie Tsai, and Jens K. Nørskov, *Automated*

- discovery and construction of surface phase diagrams using machine learning*, The Journal of Physical Chemistry Letters 7 (2016), no. 19, 3931–3935, PMID: 27558978.
- [102] Georges Valette, *Double layer on silver single crystal electrodes in contact with electrolytes having anions which are slightly specifically adsorbed: Part II. The (100) face*, J. Electroanal. Chem. Interfacial Electrochem. 138 (1982), no. 1, 37–54.
- [103] Stephen E. Weitzner and Ismaila Dabo, *Voltage effects on the stability of pd ensembles in pd–au/au(111) surface alloys*, The Journal of Chemical Physics 150 (2019), no. 4, 041715.
- [104] Stephen E. Weitzner and Ismaila Dabo, *First principles simulations of electrified interfaces in electrochemistry*, Heterogeneous Catalysts (2021).
- [105] T. M. L. Wigley, *The pre-industrial carbon dioxide level*, Climatic Change 5 (1983), no. 4, 315–320.
- [106] Kuan-Yu Yeh and Michael J Janik, *Density Functional Theory Methods for Electrocatalysis*, Computational Catalysis (Aravind Asthagiri and Michael J Janik, eds.), Royal Society of Chemistry, dec 2013, pp. 116–156.