The Abdus Salam International Centre for Theoretical Physics

PhD course in Theory and Numerical Simulation of Condensed Matter

Interatomic Potential for Li-C Systems from Cluster Expansion to Artificial Neural Network Techniques

Candidate
Yusuf Shaidu

Supervisor
Prof. Stefano de Gironcoli

A thesis presented for the degree of Doctor of Philosophy

Trieste, October 2020

List of publications

1 **Shaidu Y.**, Lot R., Pellegrini F., Kucukbenli E. , Kaxiras E. and De Gironcoli S., *Systematic Generation of Accurate Network Potentials: the Case of Carbon*, (2020) (arXiv:2011.04604)

2. PANNA: Properties from Artificial Neural Networks Architecture, `https://gitlab.com/PANNAdevs/panna`

3. Lot R., Pellegrini F., **Shaidu Y.** and Kucukbenli E. ,*PANNA: Properties from Artificial Neural Network Architectures*, Computer Phys. Comm. 256 107402, 2020.

4. **Shaidu Y.**, Kucukbenli E. and De Gironcoli S., *Lithium Adsorption on Graphene at Finite Temperature*, J. Phys. Chem. C 122 (36), 20800-20808 (2018).

Acknowledgement

I would like to thank my supervisor, Prof. Stefano de Gironcoli for his dedication to making me a better researcher. I am immensely grateful for his encouragement, care and prompt response to all my scientific and non-scientific questions and queries.

This thesis would not have been a success without the tremendous help from Emine, Franco and Ruggero. My current knowledge of machine learning techniques emanated from them and subsequent periodic discussions with them led to the tremendous achievements recorded in this thesis.

I cannot forget to say a big thank you to all my colleagues in the Condensed Matter Sector in particular and SISSA at large. My time in SISSA was enjoyable as it was all because of you. In particular, my sincere gratitude goes to my "immediate" colleagues: Andrea, Claudio, Giulia, Karla, Lorenzo, Silvia, Matteo F. and Matteo W. It was indeed a great time that cannot be forgotten in a hurry.

Special thanks to my lovely son and wife for their understanding. Despite all odds and unprecedented eventualities, my wife's persistent support and encouragement contributed to my mental well-being and hence, the success of this work. Thank you for being there for me. To all my family members, I say thank you for be there for me.

I sincerely thank ICTP for providing me with this great opportunity to pursue my PhD at SISSA. Special thanks to all the academic and non-academic staff of the center.

Finally, I would like to thank SISSA for the computer resources they provided that make the research work possible.

Contents

Chapter 1

Introduction

The environmental concern posed by the use of fossil fuel energy sources as well as the diminishing natural reserves of the fuel has motivated the adoption of clean and renewable energy sources. Renewable energy sources are needed, for instance, to replace the combustion engine vehicles with electric and/or hybrid vehicles for environmentally friendly transportation systems. To finally adopt the renewable energy technology, an efficient energy storage system must be developed. The electrochemical energy storage system is the most prominent one in this regards, and lithium ion batteries (LIBs), a technology that is currently used in household electronic devices, portable electronics and in electric vehicles, have a lot of prospects because of their light weight, high energy density and long life cycle.

The LIB is a rechargeable battery that is made up of a graphitic anode, an electrolyte and a cathode, commonly a layered oxide material that serves as Li source. During charging, Li ions migrate from the cathode and intercalate on the graphitic anode and vice-versa while discharging. The battery is fully charged when the graphitic anode reaches a saturation with respect to the Li uptake, therefore, the amount of Li ions on the anode determines the capacity of the battery. Thermodynamically, the most stable configuration of Li intercalated graphite is one Li to every six carbon which corresponds to a theoretical specific capacity of 372 mAh/g. This undoubtedly defines the limit of the achievable capacity of the LIBs. For large scale applications, high capacity LIBs are desirable and this necessitates the search for alternative anode materials.

Graphene based anodes have been suggested as possible alternative high capacity anode [1] due to the presence of large surface area accessible for Li adsorption. However, experimental investigation of Li capacity on single layer graphene [2] revealed that the capacity is less than that of graphite due to strong Coulomb repulsion between Li ions on either side of the graphene layer. Interestingly, experimental studies on non-graphitic arrangements of graphene nanoflakes have reported high Li-capacity in the 540-1500 mAh/g range [1,3–7], a capacity higher than that of the graphite anode. These

experimental observations suggest an apparent discrepancy where the building block, graphene, fails to host significant amount of Li ions while graphene-based materials can be very promising anode materials, despite a major amount of single-layer graphene presence within them. Moreover, theoretical studies of Li on pristine graphene [8] suggest that the Li-graphene system at zero temperature is thermodynamically unstable with respect to phase separation into graphene and metallic Li confirming the experimental results by Pollak *et al.* [2]. However, the effects of defects and vacancies in the carbon host have also been studied within density functional theory (DFT) [8–16] and the results suggest that although the stability of the composite systems is enhanced under these conditions, the stability fades away as the density of defects/vacancy decreases; another confirmation that pristine graphene is unstable at zero temperature.

Most of the theoretical studies mentioned above are done in ideal scenario and the systems are modelled with few atoms per cell as well as few graphene layers [8] in some cases due to the cost of performing DFT calculations of very large systems. Therefore, an extensive simulation of Li interaction with graphene materials in experimental conditions is required to understand the mechanisms that stabilize high Li density in graphene materials and the suitability of the graphene materials as an alternative anode for LIB.

In this thesis, an attempt is made to provide atomistic explanation to the experimentally observed high storage of Li^+ ions on graphene based anode materials of lithium ion batteries through extensive simulation of graphene based anode in experimental conditions. The most part of the work is devoted to methods development required for the study. The thesis is structured as follows: In Chapter 2, we start with a brief review of the general framework of density functional theory (DFT) and the approximations to the exchange correlation functional including the extension to the non-local approximation to the correlation energy that allows to describe the van der Waals interactions. In Chapter 3, we provide a brief overview of cluster expansion method, Monte Carlo and thermodynamic integration techniques as employed to study Li adsorption on graphene substrate. In this chapter, we present the results of Li interaction with graphene both at zero temperature and at finite temperatures with a bi-dimensional Li adsorption on graphene substrate. This approximation is necessitated by the limitations of the cluster expansion methods. As a result of this limitation, we adopt a more general, sophisticated and very successful methodology for interatomic potential construction, the artificial neural network techniques and the detail of the methods and implementation are discussed in Chapter 4. In Chapter 5, a systematic procedure to generating accurate neural network potentials (NNP) is proposed. A successful application of the proposed method to construct a general purpose interatomic potential for carbon is presented. Furthermore, in Chapter 6, we present an extension of the NNP scheme to include long range electrostatics through the charge equilibration

scheme, an energy contribution paramount for proper description of Li-C interaction due to charge transfer from Li to the C host material. The preliminary results obtained for Li-C interactions with and without the long-range correction are presented. Finally, we summarize our work and give possible future directions in Chapter 7.

Chapter 2

Theoretical Background

In this chapter, we give an overview of the main theoretical framework employed for the *ab-initio* studies throughout this work. We start with a brief overview of the general density functional theory (DFT) framework followed by a short description of the main approximations to the exchange and correlation energy functional in DFT and finally review the various treatments of van der Waal (vdW) interactions in DFT.

2.1 Density Functional Theory

The Hamiltonian of a many electron system we are interested to solve is given by

$$\begin{aligned} H = &- \sum_i \frac{\hbar^2}{2m}\nabla_i^2 - \sum_I \frac{\hbar^2}{2M_I}\nabla_I^2 + \frac{K}{2}\sum_{ij}\frac{e^2}{\mid \mathbf{r}_i - \mathbf{r}_j \mid} \\ &- K\sum_{iI}\frac{Z_I e^2}{\mid r_i - \mathbf{R}_I \mid} + \frac{K}{2}\sum_{IJ}\frac{Z_I Z_J e^2}{\mid \mathbf{R}_I - \mathbf{R}_J \mid} \end{aligned} \tag{2.1}$$

$$H = T_e + T_N + V_{ee} + V_{eN} + V_{NN} \tag{2.2}$$

where m and M_I are the mass of electrons and nuclei respectively, $K = \frac{1}{4\pi\epsilon_0}$, ϵ_0 is the permittivity of free space. T_e and T_N are the kinetic energy operators for the electrons and nuclei and V_{ee}, V_{eN} and V_{NN} are electron-electron, electron-nucleus and nucleus-nucleus interaction potential operators respectively.

A closed form solution of this Hamiltonian is not known for more than one electron systems. In order to describe many electron systems, an approximate solution to equation 2.1 becomes inevitable. The very fundamental approximation is based on the fact that the electrons are much lighter than nuclei. This allows the electronic and nuclear degrees of freedom to be decoupled and the many body wavefunction therefore can be written as a product of the electronic wavefunctions at fixed nuclei coordinates, $\mathbf{R}$, and the nuclear wavefunctions. This approximation is referred to as the Born-Oppenheimer approximation [17]. The approximation gives rise to two

decoupled equations namely, the electronic equation in which the nuclear coordinates enter as parameters given by

$$\left[-\sum_i \frac{\hbar^2}{2m_i}\nabla_i^2 + \frac{K}{2}\sum_{ij}\frac{e^2}{\mid \mathbf{r}_i - \mathbf{r}_j \mid} - K\sum_{iI}\frac{Z_I e^2}{\mid \mathbf{r}_i - \mathbf{R}_I \mid} + \frac{K}{2}\sum_{IJ}\frac{Z_I Z_J e^2}{\mid \mathbf{R}_I - \mathbf{R}_J \mid}\right] \times$$
$$\psi(\{\mathbf{r}_i\},\{\mathbf{R}_I\}) = E(\{\mathbf{R}_I\})\psi(\{\mathbf{r}_i\},\{\mathbf{R}_I\}), \tag{2.3}$$

and the nuclear dynamics described by

$$\left[-\sum_I \frac{\hbar^2}{2M_I}\nabla_I^2 + E(\{\mathbf{R}_I\})\right]\Phi(\{\mathbf{R}_I\}) = \varepsilon\Phi(\{\mathbf{R}_I\}), \tag{2.4}$$

where $E(\{\mathbf{R}_I\})$ is the electronic energy at fixed nuclear coordinates known as the potential energy surface (PES).

Despite this simplification, the electronic problem is still not tractable due to the presence of the electron-electron interaction term. One approach to solve this electronic problem is Density Functional Theory (DFT) formulated by Hohenberg and Kohn in 1964 [18]. The DFT is based on two theorems which state as follows:

1 *For any system of interacting electrons in an external potential $V_{ext}(r)$, there is a one-to-one correspondence between $V_{ext}(r)$ and the ground state electron density.*

2 *There exists a universal functional of the density, $F[n]$, such that for any given external potential, $V_{ext}(r)$, the total ground state energy of the electronic system, $E[n]$, is given by the minimization of the functional*

$$E[n] = F[n] + \int V_{ext}(\mathbf{r})n(\mathbf{r})d\mathbf{r}. \tag{2.5}$$

The universal functional is

$$F[n(\mathbf{r})] = \min_{\psi\to n}\langle\psi|(T + V_{ee})|\psi\rangle = T[n(\mathbf{r})] + V_{ee}[n(\mathbf{r})]$$

where $T[n(\mathbf{r})]$ is the kinetic energy functional and $V_{ee}[n(\mathbf{r})]$ the electron-electron interaction functional.

The Hohenberg and Kohn theorems imply that any ground state property of the electronic system can be expressed as a functional of the ground state density and hence can be determined once the ground state density is known. However, the prescription to construct the universal functional is not provided, which renders the Hohenberg and Kohn formulation impractical. To address this challenge, Kohn and Sham (KS) [19] proposed a mapping of the interacting many-body particle system unto an auxiliary system of non-interacting particles in an effective potential, v_{KS}, with the requirement

that the ground state density of the non-interacting system equates to the density of the interacting system, leading to a single particle eigenvalue problem. In order to determine v_{KS}, Kohn and Sham decomposed the Hohenberg-Kohn energy functional in the following form

$$E_{KS}[n] = T_s[n] + \int V_{ext}(\mathbf{r})n(\mathbf{r})d\mathbf{r} + E_H[n] + E_{xc}[n], \tag{2.6}$$

where $T_s[n]$ is the kinetic energy of the non-interacting electrons and it is written in terms of the single-particle wavefunctions as:

$$T_s[n] = -\frac{\hbar^2}{2m}\sum_{i=1}^{N}\int \psi_i^*(\mathbf{r})\nabla^2\psi_i(\mathbf{r})d\mathbf{r}. \tag{2.7}$$

$E_H[n]$ is the part of the electron-electron interaction that can be calculated as the electrostatic interaction of the density distribution, $n(\mathbf{r})$:

$$E_H[n] = \frac{e^2}{2}\int \frac{n(\mathbf{r})n(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|}d\mathbf{r}d\mathbf{r}' \tag{2.8}$$

The $E_{xc}[n]$ term is the so called exchange and correlation (XC) energy functional that contains all the unknown parts of the energy functional.

The ground state of the KS auxiliary system is obtained by minimizing the KS energy functional in equation 2.6 with respect to the single particle wavefunctions subject to the condition of fixed total number of electrons of the system. This gives rise to the KS equations that reads as follows:

$$H_{KS}\psi_i(\mathbf{r}) = \epsilon_i\psi_i(\mathbf{r}). \tag{2.9}$$

Here H_{KS} is the Kohn-Sham effective hamiltonian given by

$$H_{KS} = -\frac{\hbar^2}{2m}\nabla^2 + v_{KS}(\mathbf{r}), \tag{2.10}$$

$$v_{KS}(\mathbf{r}) = v_{ext}(\mathbf{r}) + e^2\int \frac{n(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|}d\mathbf{r}' + v_{xc}(\mathbf{r}), \tag{2.11}$$

where $v_{xc}(\mathbf{r}) = \delta E_{xc}[n]/\delta n(\mathbf{r})$ is the exchange-correlation potential.

Given the external potential, i.e. the atomic positions, the Kohn-Sham equations (equation. 2.9- 2.11) depend only on the density through the single particle wavefunctions which are the solution of the KS system of equations 2.9 itself. Therefore these equations must be solved self-consistently. The self-consistency loop starts from a guessed single particles wavefunction from which the density is computed. This is

followed by a computation of the KS potential and finally plug in the KS problem to solve for the single particle orbitals. This procedure is continued until self consistency is reached. Once the solution is found, the total energy of the interacting electrons system can be calculated in terms of the eigenvalues, ϵ_i, of the KS equations as:

$$E[n] = \sum_i \epsilon_i - E_H[n] + E_{xc}[n] - \int v_{xc}(r)n(r)dr \tag{2.12}$$

It is worthwhile to mention that the independent particle kinetic energy and the Hartree energy that constitute the majority of the energy have been explicitly factored out. Therefore, if the exact form of the exchange-correlation functional,$E_{xc}[n]$, were known, the solution of the KS system would give the exact ground state density and hence, the total energy of the fully interacting system would be exactly determined. However, the $E_{xc}[n]$ term is unknown and one has to resort to approximations. Since the $E_{xc}[n]$ is expected to contribute a small fraction to the total energy, any reasonable approximation to $E_{xc}[n]$ should give a good estimate of the ground state properties of the interacting system.

The simplest of these approximations is the local density approximation (LDA) [19], based on the properties of the homogeneous electron gas in which $E_{xc}[n]$ is defined as the integral of the energy density per electrons of the homogeneous electron gas evaluated at the density of the inhomogeneous electron system at a given point $\mathbf{r}$, i.e

$$E_{xc}^{LDA}[n] = \int n(r)\epsilon_{xc}(n(r))dr, \tag{2.13}$$

where $\epsilon_{xc}(n(r))$ is the exchange and correlation energy density per particle whose functional form [20–22] are based on accurate Quantum Monte Carlo (QMC) simulations.

Another popular approximation to the XC functional is the semi-local approximation referred to as the Generalized Gradient Approximation (GGA). This is an extension of the LDA to include local variation of the density at a fixed point r. In this case, the exchange and correlation energy functional now depends also on the gradient of the density at $\mathbf{r}$ and reads as

$$E_{xc}^{GGA}[n] = \int n(r)F_{xc}(n(r), |\nabla n(r)|)\epsilon_x(n(r))dr, \tag{2.14}$$

where $\epsilon_x(n(r))$ is the exchange energy density per particle of a homogeneous electron gas and F_{xc} is an enhancement of the exchange-correlation energy density, ϵ_{xc}. The dependence of $E_{xc}^{GGA}[n]$ on the variation of the density enters through the enhancement factor, F_{xc}. Depending on the modeling of F_{xc}, different GGA functionals exist. The most commonly used GGAs are the PBE proposed by Perdew, Burke and Ernzerhof [23] and BLYP which is a combination of the exchange energy functional expressed by Becke

in 1988 [24] with the correlation energy functional of Lee, Yang and Parr [25].

The (semi)local approximation so far does not take into account non-local electron correlations which are important to describe spatially separated fragments such as dimers, layered materials, biological micromolecules, molecular crystals, physisorption etc. This missing energy contribution due to density fluctuations are called the van der Waals (vdW) interactions. In this work, we intend to study systems with significant inhomogeneity in the charge density distribution such as Li interacting with single layer and multilayer graphene as well as pure carbon systems. This therefore necessitates the use of functionals beyond the (semi)local functionals already introduced. In the following section, we provide a short review of the development of the vdW interaction correction in DFT up to the revised Vydrov and Van Voorhis (rVV10) [26] functional that incorporates long-range vdW interactions.

2.2 van der Waals Interactions in DFT

Dispersion interactions are due to correlations between density fluctuations in two spatially separated fragments (atoms or molecules) in solids. Although the DFT formalism is exact and in principle contains all the electron correlations including the long-range vdW interactions, the widely used (semi)local approximation to the exchange-correlation functional misses these interactions. The inability of the (semi)local approximation to describe the dispersion interactions was considered in 1998 by Kohn et al. [27]. In their work, the electron-electron interaction in the many body Schroediger equations was split into short- and long-range part, and the long-range parts that contributes to the vdW energies was expressed in terms of density-density response functions which is formally derived from the adiabatic connection fluctuation dissipation formula (ACFD) [28, 29]. It was shown that for two spatially separated spherical atoms A and B, the asymptotic vdW interactions between the two atoms at separations R is $E_{vdW} = -\frac{C_6}{R^6}$, a well known form of the London dispersion; where

$$C_6 = \frac{3}{\pi}\int_0^\infty d\omega \alpha_A(\omega)\alpha_B(\omega) \tag{2.15}$$

and α_A and α_B are the dipole-dipole polarizability of the isolated atom A and B, respectively.

Nowadays, the treatments of vdWs interactions in DFT can be classified into two broad categories, namely: (i) the empirical pair-wise dispersion correction in which a model description of the long-range vdW energy, mostly R^{-6} term is added to the KS total energy, and (ii) the fully nonlocal vdW density functionals that construct a vdW energy functional of the density in a self-consistent way. In the following two subsections, we provide a brief description of these categories.

2.2.1 Empirical Pair-wise vdW Corrections

The first category is based on empirical treatments [30–33]. It involves the addition of an empirical energy term to the KS total energy. The simplest form of this vdW term is written as

$$E_{vdW} = -\frac{1}{2}\sum_{AB}\frac{C_{6,AB}}{R_{AB}^6} f_{dmp}(R_{AB}) \tag{2.16}$$

where R_{AB} is the distance between atoms A and B. The function f is a short-range damping function. Depending on the scheme used to compute the $C_{6,AB}$ coefficient and the choice of the damping function, several variant of equation 2.16 exist. The first and most widely used among them being the Grimme DFT-D$_n$ [30–32] series with increasing n, as attempts are made to extend the applicability of the methods to a variety of systems. In DFT-D1 [30], the fixed atomic C_6 coefficients are used for the homonuclear coefficients, $C_{6,AA}$ while a simple combination rule ($C_{6,AB} = \frac{2C_{6,A}C_{6,B}}{C_{6,A}+C_{6,B}}$) is used to obtain the heterogeneous coefficients. The damping function in this case was constructed to avoid the divergence at small interatomic distances and minimize the contribution of the vdW corrections at close range. In the case of DFT-D2 [31], the combination rule for the heterogeneous C_6 coefficients were modelled with a simple geometric relations as $C_{6,AB} = \sqrt{C_{6,A}C_{6,B}}$ motivated by the observation that light atoms were over-weighted in the combination rule employed in the DFT-D1 scheme. The modelled C_6 coefficients in DFT-D1 and DFT-D2 are environment independent and therefore pose transferability issues. Motivated by these shortcomings, the DFT-D3 came with several modification including the determination of the C_6 coefficients from time-dependent DFT (TD-DFT) *ab initio* methods. Other modifications include extension of equation 2.16 to include higher order multipoles (C_8 and C_{10}) which are computed recursively and depend on multipoles-type expectation values based on geometrically averaged atomic densities, a three-body term and a modified damping function that satisfies the proper limits at large separation (i.e $f_{damp}(R) \to 1$ as $R \to \infty$).

Beside the DFT-D$_n$ family, other notable empirical van der Waals correction that use the electronic structure information of the system have also been devised. Becke and Johnson [34] developed an approach in which the C_6 coefficients of isolated atoms and molecules are computed based on the instantaneous dipole moment of the exchange hole. The authors later introduced the decomposition of free-atoms quantities, namely polarizability and the C_6 coefficients, following the Hirshfeld partitioning scheme [35]. In this approach, the isolated atom density is used to define the neighborhood of each atom in the molecule and on this basis, the polarizability and the C_6 coefficients are rescaled accordingly. Tkatchenko and Scheffler [33] extended this approach by rescaling, in addition to the polarizability and the C_6 coefficients, also the vdW radii thereby making them electron density dependent. All the aforementioned vdW correction schemes are done *a posteriori* upon achieving a convergence in the KS equations.

This introduce a level of inconsistency in the KS self-consistency scheme.

2.2.2 Fully Nonlocal vdW Density Functional

The fully nonlocal vdW density functionals (vdW-DF) proposed by Langreth and coworkers [36–39] is a first principles DFT treatment of the vdW of long-ranged interactions via a nonlocal electron correlation of the exchange-correlation functional. The term "nonlocal" is used to designate the dependence of the integral kernel of the dispersion energy on two electron coordinates in an inseparable manner as compared to the (semi)local treatments of the exchange-correlation functional which depends on the electron coordinate at a single point in space. The nonlocal vdW density functional depends only on the electronic density and its derivatives. The electron correlation part of the exchange and correlation energy functional can be decomposed as follows:

$$E_c[n] = E_c^0 + E_c^{nl} \tag{2.17}$$

where E_c^0 takes care of the local correlation energy, a contribution that is well captured by the (semi)local exchange correlation functional.

The nonlocal correlation energy functional is cast into the following form:

$$E_c^{nl} = \frac{1}{2} \int d\mathbf{r} \int d\mathbf{r}' \; n(\mathbf{r}) \; \Phi(\mathbf{r}, \mathbf{r}', [n]) \; n(\mathbf{r}') \tag{2.18}$$

where $\Phi(\mathbf{r}, \mathbf{r}', [n])$ is referred to as the nonlocal vdW correlation kernel. The notation $[n]$ is used to denote the dependence of the kernel on density and its gradients at **r** and **r'**. The kernel is related to the frequency dependent polarizability at two spatially separated points and obeys the asymptotic R^{-6} behavior of the dispersion interactions between those points.

The first general nonlocal vdW functional, vdW-DF1 was proposed by Dion et al. [36] in 2004. The vdW-DF1 functional is obtained from the full frequency dependent electronic polarizability in the exact adiabatic-coupling fluctuation dissipation (ACFD) theory by expanding the exact functional in powers of $S = 1 - \epsilon^{-1}$ (ϵ is the dielectric function) up to second order and assuming that the electronic polarizability depends only on the local density at point r and its gradient at that point. The resulting kernel, $\Phi(\mathbf{r}, \mathbf{r}', [n])$ has a cumbersome expression with double spatial integrals. However, the kernel can be cast into a kernel function that depends on r and r' through auxiliary variables d and d' defined as;

$$d = |\mathbf{r} - \mathbf{r}'| q_0(\mathbf{r}) \tag{2.19}$$

$$d = |\mathbf{r} - \mathbf{r}'| q_0(\mathbf{r}') \tag{2.20}$$

where $q_0(\mathbf{r}) = -\frac{4\pi}{3e^2}\left[\epsilon_c^{LDA}(\mathbf{r}) + \epsilon_x^{LDA}(\mathbf{r})\left(1 + \gamma s^2(\mathbf{r})\right)\right]$ with $\gamma = 0.8491/9$ and ϵ_x and ϵ_c are the LDA exchange and correlation energy density respectively. $s = \frac{\nabla n}{2n^{4/3}(3\pi^2)^{1/3}}$ is the reduced density gradient. The kernel can be tabulated *a priori* on a grid of d and d' therefore making the evaluation of the kernel values and its derivatives at d and d' (i.e at $\mathbf{r}$ and $\mathbf{r}'$) very fast and efficient during the self-consistency cycle. Despite the successes of vdW-DF1 description of non-local correlations in layered materials and dimers [36, 40], a detailed study of the dependence of the functional on the choice of exchange functional [41] revealed that the revised Perdew–Burke–Ernzerhof (revPBE) [42] used in the vdW-DF1 systematically underestimate hydrogen bonds. It is also well known to overestimate equilibrium separations [37, 43]. To address the aforementioned shortcomings, the vdW-DF2 [44] is proposed in which the revPBE is replaced with Perdew-Wang (PW86) [45] and the gradient correction to the LDA exchange is enhanced through the modification of γ in q_0. These modifications lead to satisfactory description of the equilibrium separations and alleviate the underestimation of the hydrogen-bond strength.

Furthermore, in 2010 [46] Vydrov and Van Voorhis proposed a very simple analytic functional form for the nonlocal correlation kernel based on their previous works [43, 47–49]. The nonlocal vdW energy functional referred to as the VV10 functional and reads as follows;

$$\Phi^{VV10}(\mathbf{r}, \mathbf{r}') = -\frac{3e^4}{2m^2}\frac{1}{\mathrm{gg}'(\mathrm{g} + \mathrm{g}')} \tag{2.21}$$

where

$$\mathrm{g} = \omega_0(\mathbf{r})\mathrm{R}^2 + \kappa(\mathbf{r}) \tag{2.22}$$

$$\mathrm{g}' = \omega_0(\mathbf{r}')\mathrm{R}^2 + \kappa(\mathbf{r}') \tag{2.23}$$

with $R = |\mathbf{r} - \mathbf{r}'|$. The other quantities are given by:

$$\omega_0(\mathbf{r}) = \sqrt{\omega_{\mathrm{g}}^2(\mathbf{r}) + \frac{\omega_{\mathrm{p}}^2(\mathbf{r})}{3}} \tag{2.24}$$

where $\omega_{\mathrm{p}}^2 = 4\pi \mathrm{ne}^2/\mathrm{m}$ is the local plasma frequency and the local band gap is defined as

$$\omega_{\mathrm{g}}^2(\mathbf{r}) = \mathrm{C}\frac{\hbar^2}{\mathrm{m}^2}\left|\frac{\nabla \mathrm{n}(\mathbf{r})}{\mathrm{n}(\mathbf{r})}\right|^4 \tag{2.25}$$

where C is an adjustable parameter to reproduce the C_6 coefficients whose optimal value is $C = 0.0093$ and

$$\kappa(\mathbf{r}) = \mathrm{b}\frac{\mathrm{v}_{\mathrm{F}}^2(\mathbf{r})}{\omega_{\mathrm{p}}(\mathbf{r})} \tag{2.26}$$

where $\mathrm{v_F} = (3\pi \mathrm{n})^{1/3}\frac{\hbar}{\mathrm{m}}$ is the local Fermi velocity and b is another adjustable parameter that controls the divergence of the R^{-6} at short range. Its optimal value is $b = 5.9$.

In order to ensure a vanishing long-range correlation energy in the limit of uniform electron density, the VV10 correlation energy functional is redefined as

$$E_c^{VV10} = E_{nl}^c + \beta N_{el} \tag{2.27}$$

where $\beta = \frac{\mathrm{e}^2}{32\mathrm{a}_0}\left[\frac{3}{\mathrm{b}^2}\right]^{2/4}$ is the negative of the energy density per electron evaluated with the uniform density limit of $\Phi^{VV10}(\mathbf{r},\mathbf{r}')$ with $\mathrm{a}_0 = \hbar^2/\mathrm{me}^2$ and N_{el} the number of electrons.

The evaluation of equation (2.18) with the kernel defined in equation (2.21) is not computationally tractable in the plane-wave framework in general. A similar problem existed in the original vdW-DF and was addressed with the introduction of a very efficient integration scheme by Roman-Perez and Soler(RPS) [50] based on the fact that the vdW-DF functional depends on the density and its gradients at $\mathbf{r}$ and $\mathbf{r}'$ through an auxiliary function $d(\mathbf{r}) = d(n(\mathbf{r}), |\nabla n(\mathbf{r})|)$ and hence allows the interpolation of the kernel on a 2-dimensional grid of points. The VV10 kernel depends on the density and its gradients in $\mathbf{r}$ and $\mathbf{r}'$ separately and hence does not allow the direct application of the RPS procedure. Sabatini et al. [26] cast the VV10 functional into a form compatible with the RPS scheme by rearranging the kernel and setting the ratio $\kappa(\mathbf{r})/\kappa(\mathbf{r}')$ as well as its inverse to unity upon detailed analysis of the ratio. The resulting functional is referred to as the revised VV10 (designated as rVV10). The rVV10 kernel reads as follows;

$$\Phi^{rVV10}(\mathbf{r},\mathbf{r}') = -\frac{3e^4}{2m^2}\frac{1}{(\mathrm{qR}^2+1)(\mathrm{q}'\mathrm{R}^2+1)(\mathrm{qR}^2+\mathrm{q}'\mathrm{R}^2+2)} \tag{2.28}$$

where $\mathrm{q}(\mathbf{r}) = \omega(\mathrm{n}(\mathbf{r}), |\nabla \mathrm{n}(\mathbf{r})|)/\kappa(\mathrm{n}(\mathbf{r}))$ and similarly for $\mathbf{r}'$ with some function of κ removed from the kernel (see the original Ref. [26] for details). The RPS interpolation scheme can now be applied directly to the rVV10 kernel in reciprocal space.

The vdW-DF functional has been shown to be sensitive to the choice of the local exchange and correlation part of the functional [41] and the revPBE functional was replaced with PW86 since it describes properly the repulsive part of vdW-DF energy. In the case of the VV10 (and also rVV10), the total exchange correlation functional comprises of the refitted Perdew-Wang for the exchange functional and the bare PBE for the correlation and reads:

$$E_{xc} = E_x^{rPW86} + E_c^{PBE} + E^{rVV10} \tag{2.29}$$

The DFT introduced in this chapter along side the exchange correlation functionals are implemented in the Quantum ESPRESSO package [51, 52], the software employed for the first principles studies throughout the thesis. Except for the purpose of comparison with standard exchange correlation functionals such as the GGA-PBE, all first

principles calculations are performed using the rVV10 functional as implemented.

In summary, we have introduced the general DFT framework and the local and semilocal approximation to the exchange-correlation functionals and provided somewhat detailed review of the two main classes of the nonlocal van der Waals treatments in the DFT framework.

Chapter 3

Lithium Adsorption on Graphene at Finite Temperature

In this chapter, we present our first principle results on Li interaction with graphene substrate accounting for the effect of temperatures. The chapter is divided into two parts. In the first part, we report the computational details and results of Li interaction with graphene at zero temperature and conclude the part with a brief motivation on the need to study the effects of temperatures. In the second part, a brief description of the cluster expansion methods, the grand canonical techniques, the thermodynamics integration techniques, the generation of fitting data and the fitting procedure are presented. Then we proceed with the discussion of finite temperature results and then conclude the chapter with a summary of the main results. See Ref. [53] for the journal publication of the work. In the subsequent chapters, the attempts we have made to address the limitations of the cluster expansion methodology are presented.

3.1 Introduction

The need for high energy-density Lithium ion battery (LIB) technology is increasing with the increasing demand for portable electronic devices and electric vehicles. While alternatives are being developed, graphite has long been the typical negative electrode used in LIBs [54–56]. During charging, the Li ions intercalate between the layers of graphite which, at ambient conditions the most stable phase, reaches a Li to C ratio of 1:6 [57], corresponding to a theoretical specific capacity of 372 mAh/g. Graphene, the single layer of carbon atoms from graphite [58], has also been suggested as potential Li host by itself [1] due to the presence of large surface area accessible for Li adsorption. However, an experimental work by Pollak *et al.* [2] revealed that Li capacity of a single-layer graphene is less than that of graphite due to strong Coulomb repulsion between Li ions on either side of the graphene layer. Nevertheless, experimental studies on

non-graphitic arrangements of large graphene flakes have reported high Li-capacity in the 540-1500 mAh/g range [1, 3–7], surpassing the graphitic anode capacity. These experimental observations suggest an apparent discrepancy where the building block, graphene, fails to host significant amount of Li ions while graphene-based materials can be very promising anode materials, despite a major amount of single-layer graphene presence within them.

Several theoretical studies based on Density Functional Theory (DFT) [8–16] attempted to elucidate the nature of high Li adsorption on graphene-based systems: Lee and Persson [8] showed that Li cannot reside on the surface of a defect-free single layer graphene in equilibrium with bulk Li metal and Li capacity would be lower than that of graphite. Li adsorption on graphene with defects, vacancies [9, 10] and edges [11–13] has also been investigated and the results revealed that the binding energy of Li decreases rapidly as its position moves away from the vicinity of the edges and/or defects, meaning that defects alone would provide a relatively small, local, enhancement to the adsorption capacity. Hydrogen passivation of defects results in positive binding energy [10] with respect to bulk Li, similar to that of defect-free graphene, making the capacity of hydrogen-passivated defective graphene inferior to that of graphite. Therefore these theoretical studies do not provide a satisfactory explanation for the experimental observation of high Li uptake of graphene-based systems. It is important to note that these works explore single, dilute Li ion adsorption configurations. However, studies on compact Li clusters show that the energetic stability of Li clusters on graphene increases with the cluster size [14–16] suggesting that the formation of compact Li clusters could explain the reported high Li uptake by graphene materials. However, the possibility of high Li uptake due to formation of Li clusters is not necessarily desired since a continuous growth of clusters on the anode may result in reduced battery efficiency and compromise safety of the battery. Therefore the desired mechanism for high Li uptake is a dispersed configuration rather than clustering, especially if cluster growth is thermodynamically favored. Noting that the theoretical calculations mentioned so far were done at zero temperature while LIBs operate at room temperature or above, it is important to investigate whether thermal effects may favor dispersed Li configurations or may stabilize small cluster sizes.

Thermal effects on the adsorption of Li clusters on graphene were previously estimated in Refs. [15] and [16] by choosing a random configuration as a reference, approximating the entropy through it, and by neglecting the configurational entropy of the clusters. In these studies, the formation of Li clusters was found to be stable against random distribution of Li ions for high concentration (1Li:6C) at 300 and 500 K; while at low concentrations (1Li:72C and lower) a disperse configuration was favored. Furthermore, Liu *et al.* [15] estimated nucleation barriers as high as approximately 15 eV for a cluster to form in dilute configurations (1Li:162C), suggesting that thermal effects

may indeed stabilize disperse configurations. However these findings rest on a heavily approximated description of entropy, which is particularly critical in understanding the thermal effects.

The rest of the chapter is organized as follows: in Sec. 3.2.1, the computational details used for the zero temperature studies are presented and the results are presented in Sec. 3.2.3. The methods employed for finite temperature studies, namely: the cluster expansion methods, the fitting dataset, the fitting procedure, the Monte Carlo techniques and the thermodynamic integration, are presented in Sec. 3.3.1, 3.3.2, 3.3.3, 3.3.4 and 3.3.5 respectively. The finite temperature results are discussed in 3.3.6, 3.3.7 and 3.3.8. Finally, we present our conclusions and possible future directions.

3.2 Zero Temperature Calculations

In this section, we start with the presentation of the methodology used for the zero temperature studies and then proceed with the discussion of the results obtained at zero temperature and motivate the need for finite temperature studies.

3.2.1 Computational Details

We examined Li adsorption on graphene at zero temperature by means of Density Functional Theory (DFT) [18, 19] within the planewave-Projector Augmented Wave (PAW) [59] framework using the Quantum ESPRESSO package [51, 52]. The PAW pseudopotentials were generated from the pslibrary.1.0.0 [60]. Kinetic-energy cutoffs of 125 Ry and 500 Ry were used in the expansion of wavefunctions and charge density respectively. With these parameters, a total energy convergence of 0.1 mRy/atom was achieved. The revised Vydrov and van Voorhis (rVV10) [26, 46] exchange-correlation functional was used to account for non-local van der Waals interactions. For comparison, we also performed calculations using generalized gradient approximation in the Perdew-Burke-Ernzerhof (PBE) parametrization [23]. The total energy calculations of Li-graphene were done in graphene supercells with dimensions ranging from 2x2 to 8x8 primitive cells. The brillouin zone was sampled uniformly according to the Monkhorst-Pack [61] procedure with Methfessel-Paxton [62] smearing of 0.02 Ry. A 30x30x1 mesh centered around gamma point was found to be sufficient for 0.1 mRy per atom convergence of total energy of graphene primitive unit cell. Calculations with different supercells were done conserving this sampling density.

The in-plane lattice parameters of pristine graphene were fully optimized while the spacing between periodic replicas along the perpendicular direction was fixed at 20 Å, a distance at which the effect of periodic replicas are shown to be negligible (See Sec. 3.2.2 for details). In each configuration, all internal coordinates were relaxed until force

components were reduced below 0.1 mRy/Bohr. The optimized lattice parameter of graphene with rVV10 functional is 2.470 Å (2.466 Å with PBE), in good agreement with previous DFT calculations of 2.463 Å [10], 2.45 Å [11], 2.47 Å [63], 2.46 Å [64], as well as the experimental in-plane lattice parameter of graphite 2.459 Å [65].

The stability of Li upon adsorption on graphene was investigated by measuring the adsorption energy per Li ion, E_{ad}, with respect to bulk Li:

$$E_{ad}(n) = \frac{E(n) - E(n=0) - nE_{Li}}{n} \tag{3.1}$$

where $E(n)$ is the calculated total energy of the Lithium-adsorbed graphene of a given supercell size, $E(n = 0)$ is the pristine graphene with the same supercell size and E_{Li} is the energy per atom in bulk Li in the body centred cubic (bcc) phase. When indicated as such, we also report adsorption energy with respect to the isolated Li atom in vacuum for comparison with literature. The energy difference between Li in vacuum and bulk is the cohesive energy which is calculated as 1.586 eV for rVV10 and 1.561 eV for PBE.

3.2.2 Effect of Periodic Boundary Conditions

The first principle calculations of Li on graphene were performed at fixed vacuum size between adjacent plane of graphene of 20 Å as mentioned in Sec. 3.2.1. To demonstrate the impact of this parameter on the adsorption energy, we perform Li adatom adsorption calculations on a 3x3 graphene supercell while varying the vacuum size, i.e. graphene layer separation (LS) from 10 to 30 Å in a step of 2 Å. The single Li adatom on graphene case has the highest amount of charge transfer, therefore, it is expected to have the most pronounced long-range electrostatic behaviour between graphene sheets. Hence the error reported here can be seen as an upper bound to the impact of periodic boundary condition (PBC). This analysis is needed to estimate the error incurred in making an artificial periodicity along the perpendicular direction to the graphene plane.

Fig. 3.1 shows the adsorption energy, E_{ad}, referenced to bulk Li as a function of layer separation (LS). The adsorption energy changes by more than 50 meV between LS=10 and LS=20 Å and only less than 10 meV between LS=20 and 30 Å indicating that the adsorption energy at LS=20 Å is considerably close to its converged value. In order to estimate the adsorption energy at infinite separation, we fit the DFT data (red curve of Fig. 3.1) to a simple function of LS and the in-plane distance between Li adatoms on surface , R:

$$E(LS) = E_{ad}(\infty) + \frac{b}{LS} + \frac{c}{\sqrt{(R^2 + LS^2)}} \tag{3.2}$$

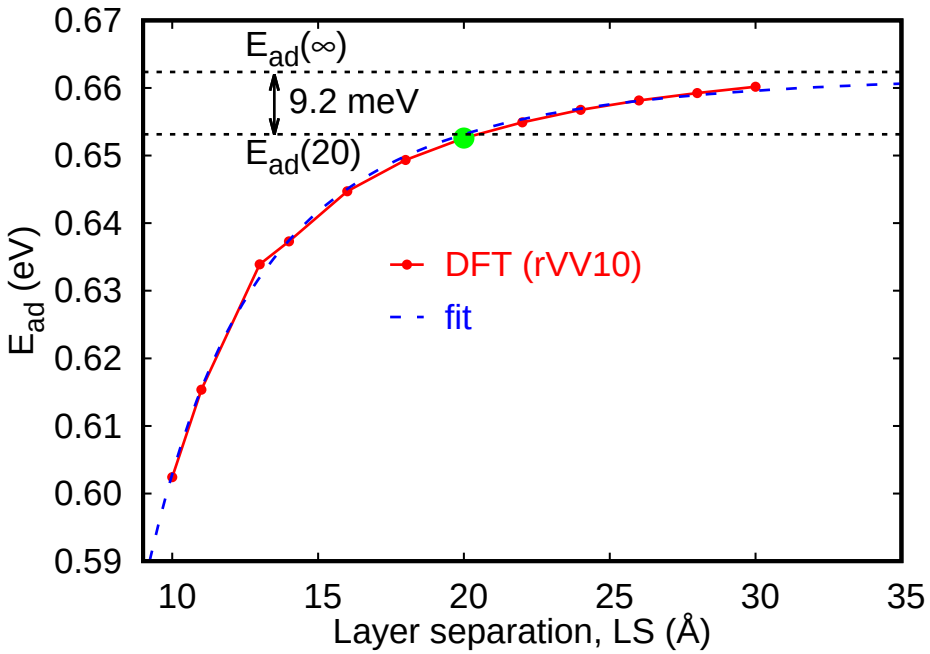

Figure 3.1: Adsorption energy(E_{ad}) of Li adatom on the hollow site of 3x3 graphene supercell as a function of layer separation (LS). The infinite layer separation limit, $E_{ad}(\infty)$ was obtained via a fit using Eq.3.2. The parameters of the fit are $E_{ad}(\infty)$ = 0.662362, b=-3.04981;c=3.05586. R is the nearest neighbor Li-Li distance on the plane of graphene, R=3×2.47 Å in the case of 3x3 supercell. The root-mean-square (rms) error of the fit is 0.77 meV.

This functional form mimics the Coulomb interaction of Li adatom and its periodic images on adjacent graphene layers, as well as the interaction with the periodic images of its nearest neighbors. Therefore at infinite layer separation, we can approximate the energy as $E_{ad}(\infty)$. The difference in adsorption energy between LS=∞ and LS=20 Å is estimated to be 9.2 meV. This value is much smaller than the energy difference between Li-adatom and Li-cluster phases at 0 K (which ranges between 44 and 176 meV on a 3x3 supercell and between 149 and 312 meV on a 6x6 supercell), therefore the vacuum size of 20 Å is sufficient for an accurate description of the stability of Li adsorption on graphene at 0 K as discussed in the next section. The effects of the PBC on the finite temperatures energetics will be discussed later.

For comparison, here we mention a few of the layer separation values used in previous works, namely, Ref. [8], Ref. [15] , Ref. [10], Ref. [9] and Ref. [66] used layer separations of 15, 15, 14.78, 12 and 20 Å respectively. In these calculations, unlike ours, the non-local van der Waals interactions (which are effective in longer ranges than standard semilocal functionals) were not included. This may explain their choice of smaller LS. Indeed, in Ref. [14] and Ref. [67] the van der Waals interactions were taken into account, and layer separations of 18 and 20 Å were used respectively.

Similarly, the planar average of the charge density difference in Fig.3.6 shows that the charge density profile does not change significantly at layer separations above 16 Å leading to the same conclusion as in the adsorption energy analysis.

3.2.3 Energetic Stability of Li Adsorption at Zero Temperature

Here, we analyze the energetic stability of the three main binding sites of graphene, namely the top (T), the bridge (B) and the hollow (H) sites (see Fig. 3.2) and the results are shown in Table 3.1 and also depicted in Fig. 3.3 (red line). The results in

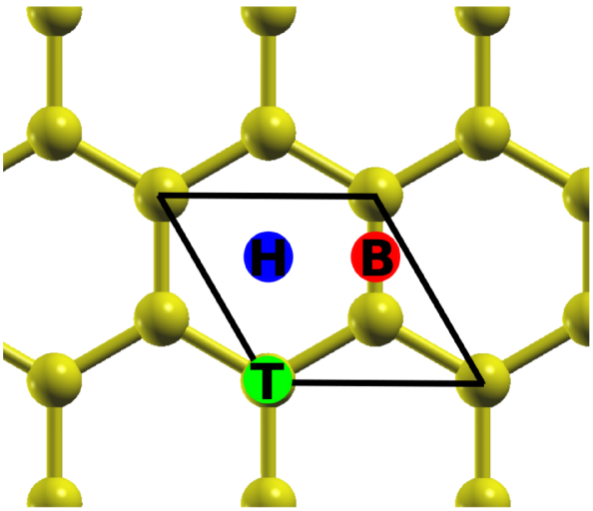

Figure 3.2: Lithium binding sites on graphene are commonly named as hollow (H), top (T) and bridge (B) sites. The black lines are used to map the primitive unit cell of graphene.

Table 3.1 show that Li atom binds preferably to the H site, in agreement with previous DFT calculations [9,11,63,66]. The energy difference, ΔE_{ad} between H site and that of B and T is beyond a few multiples of $k_B T$ at room temperature, and gets significantly higher at lower concentrations. Therefore at concentrations corresponding to 1Li:8C (single Li atom in a 2x2 supercell) and lower, Li atoms would be expected to occupy mainly the H site. These observations allow the adsorption of Li on graphene, to a good approximation, be restricted to the H-site only. To this end, we limit our investigations of Li adsorption on graphene to the H site in the rest of the chapter.

We then investigate the energetic stability of small Li-clusters as a function of supercell size (see Fig. 3.3). For all the concentrations considered, the adsorption energy calculated with respect to bulk Li is positive, which suggests that Li may not reside on the surface of a defect-free single layer graphene in equilibrium with Li-metal, in agreement with previous DFT calculations [8]. The adsorption energy per Li decreases with increasing cluster size for a fixed cell size. Similarly, at fixed cluster size, the adsorption energy decreases with increasing inter-cluster distance. The adsorption energy further decreases going from planer Li adsorption to 3D adsorption which explains the stability of the Li bulk phase.

A qualitative difference between single adatom and cluster behavior can be observed. For the single Li adatom case, the general trend of the adsorption energy at different binding sites is the same for rVV10 and PBE functionals (see Table 3.1). Similarly, both functionals capture the correlation with cell size qualitatively the same

Table 3.1: The adsorption energy of Li at the hollow site of graphene, E_{ad}, calculated with respect to bulk Li in bcc phase (values referenced to isolated Li atom are given in parenthesis). ΔE_{ad} is the adsorption energy difference for bridge and top sites with respect to the hollow site. All energy values are in eV.

cell size	E_{ad}	$\Delta E_{ad}(B)$	$\Delta E_{ad}(T)$
PBE			
2x2	0.712 (-0.849)	0.159	0.167
3x3	0.646 (-0.915)	0.310	0.327
4x4	0.468 (-1.093)	0.317	0.339
6x6	0.301 (-1.260)	0.294	0.316
7x7	0.270 (-1.291)	0.286	0.310
9x9	0.227 (-1.334)	0.283	0.305
rVV10			
2x2	0.625 (-0.961)	0.134	0.142
3x3	0.652 (-0.934)	0.241	0.256
4x4	0.485 (-1.101)	0.294	0.314
6x6	0.328 (-1.258)	0.267	0.288
7x7	0.290 (-1.296)	0.268	0.289
9x9	0.253 (-1.333)	0.257	0.280
Literature			
PBE [66]			
2x2	- (-0.925)	-	0.177
3x3	- (-1.025)	-	0.341
PBE [63]			
4x4	- (-1.096)	0.322	0.342
PBE [9]			
6x6	0.315 (-1.290)	0.285	0.308
7x7	0.298 (-1.308)	0.301	0.320
9x9	0.190 (-1.416)	0.356	0.376
LSDA [68]			
4x4	- (-1.567)	0.327	0.376

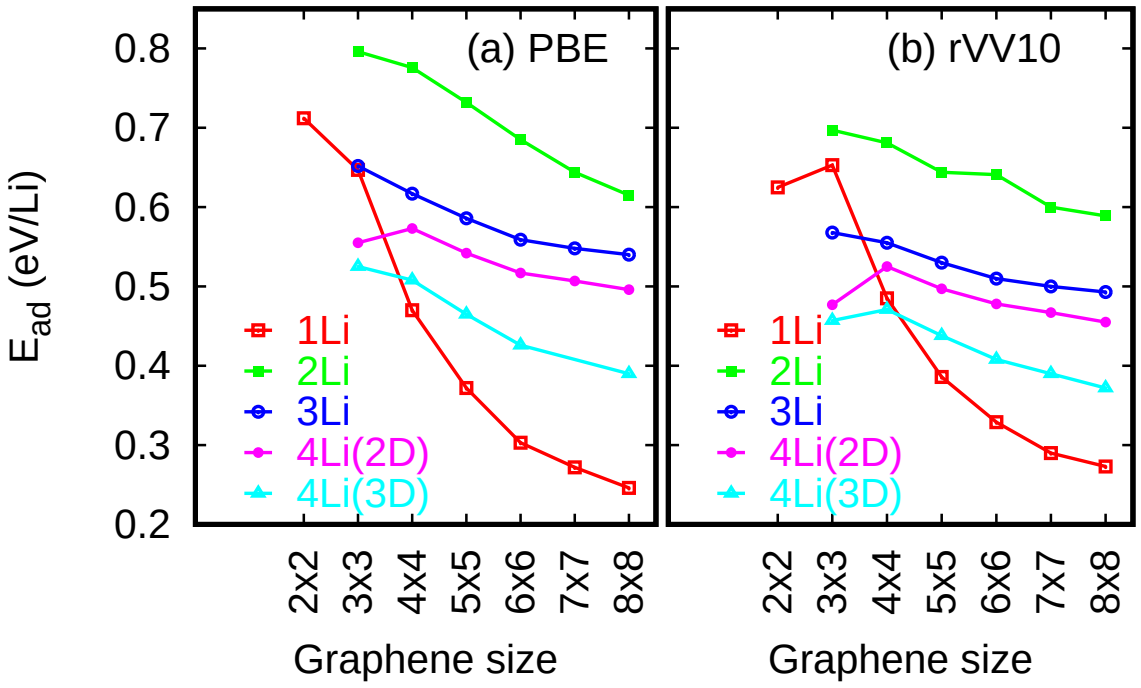

Figure 3.3: Adsorption energy of Li on graphene as a function of graphene cell size with (a) PBE and (b) rVV10 functionals, referenced to bulk Li in the bcc phase. 1Li is the Li adatom on the corresponding supercell. 2Li, 3Li, 4Li(2D) and 4Li(3D) represent Li clusters, namely Li-dimer, Li-trimer, planar Li-tetramer and pyramidal Li-tetramer respectively.

way. The most significant difference between functionals is observed in the Li cluster configurations where rVV10 functional results in significantly lower (less positive) binding energies. Comparing our results with the single adatom binding energy analysis of var der Waals (vdW) functionals in Ref. [67], where a variation of about 100 meV/Li was observed depending on the vdW treatment, it can be argued that details of how vdW interaction is treated are important for Li-graphene systems at all length scales. This is likely a combined effect of the use of different vdW schemes *and* the practice of adjusting the overall exchange and correlation description when using vdW-aware functionals.

When fixed concentration is considered, the lower binding energy with rVV10 favors the formation of Li clusters over homogeneous distribution of Li adatoms. For instance, the adsorption energy of Li adatom in a 2x2, 3x3, 4x4 supercells are higher than that of Li-clusters with four atoms in 4x4, 6x6, 8x8 supercells respectively. Same tendency is observed with PBE, although less consistently. In the rest of the discussion, we refer to the results obtained with rVV10 functional.

The observation of decreasing adsorption energy with increasing cell size in Fig. 3.3 implies the presence of electrostatic repulsion between adsorbates in adjacent simulation cells. This repulsive interaction is more evident for the single Li adatom case indicating that the dipole per atom due to charge transfer in Li-graphene system is

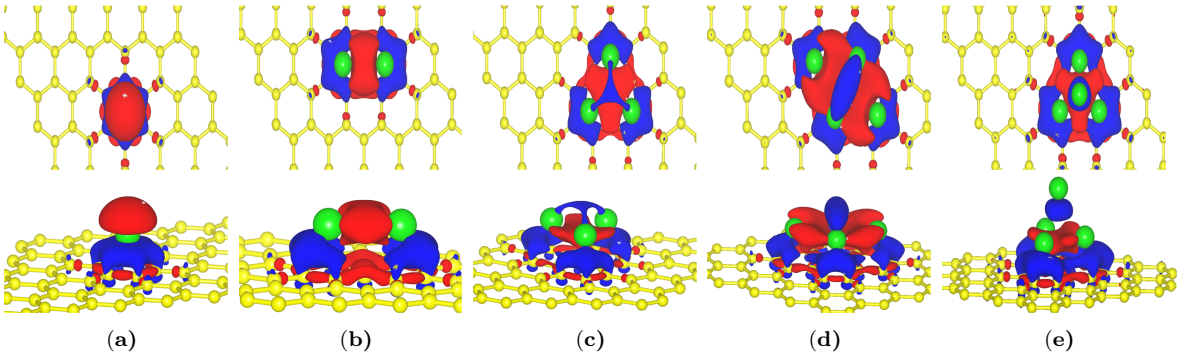

Figure 3.4: Charge density difference with rVV10 functional upon adsorption of (a) Li adatom (b) Li-dimer (c) Li-trimer (d) Li-tetramer (planar) and (e) Li-tetramer (pyramidal) on a 6x6 graphene supercell. Yellow and green spheres represent carbon and lithium atoms respectively. Red and blue regions indicate respectively the electronic charge depletion and accumulation. The surfaces are plotted at an isovalue of 0.002 $e/(a.u)^3$.

more pronounced for this configuration.

In order to gain further insight, we compute the charge density difference due to lithium adsorption:

$$\Delta\rho = \rho(graphene + nLi) - \rho(graphene) - \rho(nLi) \tag{3.3}$$

where $\rho(graphene{+}nLi)$, $\rho(graphene)$ and $\rho(nLi)$ are the charge density of Li-graphene composite, graphene supercell of equal size as in the composite and Li-cluster in the same optimized position as in the composite system respectively, and n is the cluster size. The results for different Li configurations on a 6x6 graphene supercell in Fig. 3.4 show significant charge transfer between the Li atoms and graphene. In the case of Li adatom, the electrons transferred to graphene are localized on the carbon hexagon upon which the Li atom is adsorbed, creating a net dipole (Fig. 3.4a). In the case of Li clusters (see Fig. 3.4b-e), the charge distribution is significantly different from that of Li-adatom since the positive charge is delocalized over the whole cluster, and not as elevated above the Li level as in the case of Li adatom. We quantify this difference by analyzing the planar average of the charge density difference per Li, $\Delta\rho_p(z)$, as a function of the perpendicular direction as shown in figure 3.5. A notably higher $\Delta\rho_p(z)$ is observed for Li adatom than Li-clusters. The dipole moment per Li is -0.802 eÅ for Li adatom and -0.247, -0.077, -0.125, -0.119 and -0.096 eÅ for 2, 3, 4, 5 and 6 Li-clusters respectively. This observation indicates a binary characterization of Li adsorption as single adatom versus cluster formation categories, which is utilized in the construction of the interaction potential for Li-graphene system as described in 3.3.1.

The adsorption energy reported in Fig. 3.3 also shows that isolated Li adatoms are more energetically stable than small Li-clusters at zero temperature and low concentration. This suggests that entropy may further enhance the stability of Li-adatoms over compact Li-clusters, offering the disperse Li configuration as a comparatively stable

alternative to Li clusters in the case of batteries. To this end, it becomes imperative to understand the effects of temperature on the Li-graphene system.

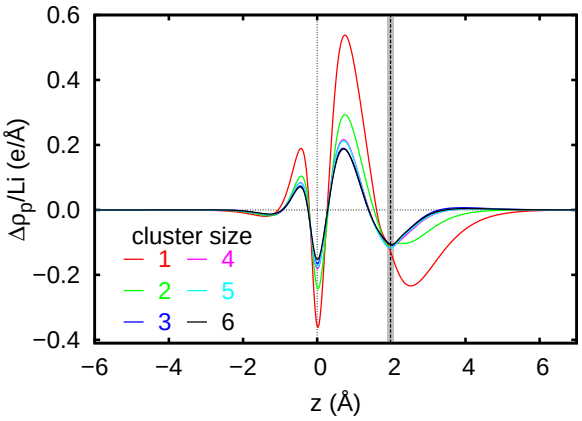

Figure 3.5: The planar average of the charge density difference per Li, $\Delta\rho_p(z) = \int \Delta\rho(x, y, z)dxdy$, for various 2D clusters as a function of the perpendicular distance from the graphene plane (situated at $z = 0\,\text{Å}$). The gray region shows the standard deviation of Li-heights around their average (the thick dashed line) above graphene plane.

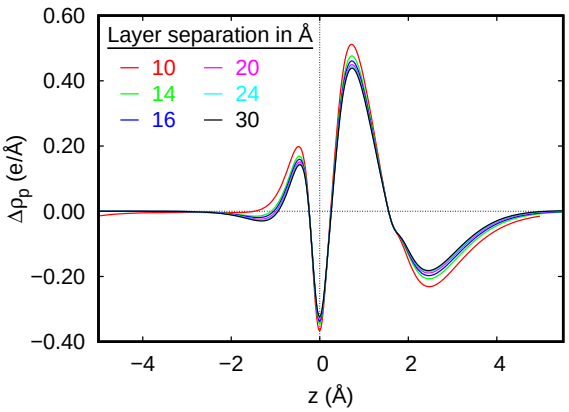

Figure 3.6: The planar average of the charge density difference between pristine graphene and Li-adsorpbed graphene, $\Delta\rho_p(z) = \int \Delta\rho(x, y, z)dxdy$, plotted at different graphene layer separations as a function of the perpendicular distance from the graphene plane (situated at $z = 0\,\text{Å}$)

3.3 Finite Temperatures Calculations

The zero temperature results reported above (see Fig. 3.3) revealed that the Li adsorption at low concentrations are energetically more stable than Li-cluster. This is a suggestion that configurational entropy which is more pronounced at low concentration than higher one is expected to further enhance the stability of Li adsorption at low concentration as the temperature effects are switched on. The behaviour of Li adsorption at finite temperature is particularly important since the LIBs with which the graphene anode is to be used, operates at least, at room temperature.

The accurate method of choice to simulate condensed matter systems at finite temperatures is the one developed by Car and Parrinello [69] in which the electrons dynamics is evolved via a fictitious equation along side the classical Newton equation for the nuclei whose driving force corresponds to the negative derivative of the total energy of the electronic problem introduced in chapter 2. This method is very expensive for large systems and requires ample amount of computer/human resources to sample the phase space of a system composed of few hundred atoms. The other extreme is the empirical force-fields which are inexpensive and therefore allow long time molecular dynamics or long time step Monte Carlo study to be carried out. However, the results obtained with empirical force-fields are, most time, only qualitatively correct with respect to accurate quantum mechanical results.

Here, we intend to construct a lithium-graphene interatomic potential, suitable to carry out an exhaustive sampling of the phase space at finite temperature, that is as affordable as the empirical force-fields and as accurate as the first principle methods. For this purpose, we adopt a site based methods called cluster expansion method as described below. Cluster expansion potentials are particularly suitable to study adsorption problem within the framework of the Grand Canonical Monte Carlo (GCMC) techniques based on Metropolis algorithm [70].

3.3.1 Cluster Expansion Techniques

The configurational energy was constructed using cluster expansion (CE) method: The Li adsorption on graphene is mapped into a lattice model that describes the energy of a given configuration of Li-graphene system as a generalized Ising-like interaction energy. To simplify the model, we restrict Li adsorption geometry to 2D and the adsorption site to the hollow site (the hexagon centers), H, of graphene. Occupation of other sites (the top of a carbon, T and the bridge, B, sites) and 3D Li configurations are beyond the scope of this chapter. Such a 2D model is compatible with experimental observations in which limited spacing between graphene sheets is found to prevent compact 3D Li growth [1, 3, 71]. In chapter 6, we present our preliminary work on the construction of a general lithium-carbon interaction potential based on artificial neural

network techniques.

In the present model, a Li occupation number, c_i, is assigned to each site i ($c_i = 0$ if site i is empty and $c_i = 1$ if site i is occupied by a Li atom, equivalent to the more common convention where spin-like variables $\sigma_i = \pm 1$ are used) and a configuration is defined as $C_N = \{c_i\}$, with $i = 1, ..., N$ for N lattice sites. The configurational energy E_{CE} is then written as a linear combination of correlation functions $\Pi_f(C_N)$ (a complete set of basis functions. see Appendix A) with each function describing a figure f (a cluster of lattice points), the corresponding coefficient in the linear expansion J_f being called the effective cluster interaction (ECI) [72]:

$$E_{CE}(C_N) = \sum_f J_f \Pi_f(C_N) = \sum_f J_f \prod_{i \in f} c_i. \tag{3.4}$$

In a lattice with N sites, there are 2^N figures: N 1-body figures, $\binom{N}{2}$ 2-body figures, $\binom{N}{3}$ 3-body figures etc. Therefore, the expansion in Eq. 3.4 has 2^N ECI's, one for each figure, requiring in principle 2^N total energy calculations to determine all the ECI's and construct an accurate CE model. Exploiting crystal symmetries (translational and rotational) and physical arguments of locality of the interaction (short-distance compact figures are expected to be more important than more-distant or dispersed ones) the number of figures to be included in the energy expansion can be constrained, leading to a limited number of independent parameters, thus requiring only a small number of total energy calculations as reference.

In this work, we include the compact figures up to 4-body (see Fig. 3.7): on-site term (figure 1), 2-body first, second and third nearest-neighbor terms up to a $2a_0$ distance, where a_0 is the lattice constant of graphene, (figures 2a, 2b and 2c), a 3-body and a 4-body term (figures 3 and 4, respectively). In addition to these, for two body interactions beyond $2a_0$, we include an effective electrostatic dipole-dipole term in the energy description. This term is set to decay with the third power of Li-Li distance, truncated for simplicity at $20a_0$. The final expansion for the total configuration energy is written as follows:

$$\begin{aligned} E_{CE}(C) &= J_1 \sum_{i \in f_1} c_i + \frac{1}{2} \sum_{p=a,b,c} J_{2p} \sum_{\{ij\} \in f_{2p}} c_i c_j \\ &+ \frac{J_3}{6} \sum_{\{ijk\} \in f_3} c_i c_j c_k + \frac{J_4}{24} \sum_{\{ijkl\} \in f_4} c_i c_j c_k c_l \\ &+ \frac{J_{dd}}{2} \sum_{ij, |\mathbf{r}_i - \mathbf{r}_j| > 2a_0} \frac{\alpha_i \alpha_j}{|\mathbf{r}_i - \mathbf{r}_j|^3} c_i c_j \end{aligned} \tag{3.5}$$

We use two parameters, J_{dd} and α, to model the effective electrostatic interaction figures. While J_{dd} is the same for all sites, α parameter is used to distinguish the

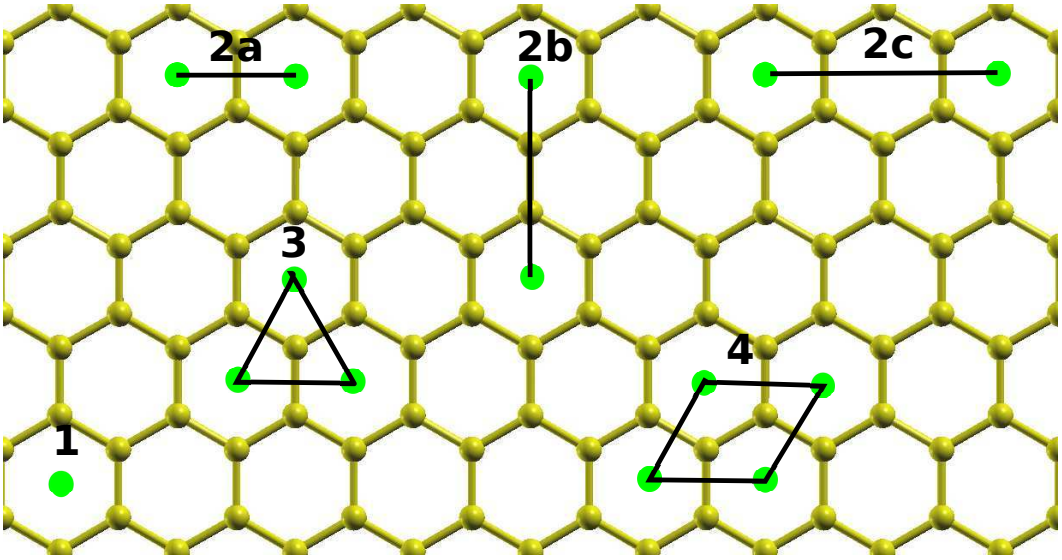

Figure 3.7: Figures included in the cluster expansion: (1) the on-site, (2) the 2-body figures between (2a) first nearest neighbors, (2b) second nearest neighbors, and (2c) third nearest neighbors, (3) the most compact 3-body figure, (4) the most compact 4-body figure.

electrostatic interaction strength in the case of the isolated atom and the clusters:

$$\alpha_i = \begin{cases} 1 & \text{if } NN_i = 0 \\ \gamma & \text{if } NN_i > 0 \end{cases} \tag{3.6}$$

where γ is an adimensional scalar parameter and NN_i is the number of occupied nearest neighbors to site i. This electrostatic description of the model is motivated by the electrostatic behavior of Li-graphene interactions analyzed in Sec. 3.2.3.

3.3.2 Fitting Dataset

The interaction energy given by equation 3.5 is fitted to DFT adsorption energies to obtain all the eight parameters of the model. A total of 25 fully relaxed configurations were used for the fitting. These configurations cover a variety of Li ion environments from single Li adatom to small cluster cases. The configurations are Li-adatom in 2x2 to 8x8 supercell, Li-dimer in 3x3 to 7x7 supercells, Li-trimer in 3x3 to 7x7 supercells and Li-tetramer in 3x3 to 6x6 supercells, the 5- and 6-Li atom clusters in the 4x4 supercell of graphene and those shown in Fig. 3.8 which are mixture of Li-clusters and adatoms.

3.3.3 Fitting Procedure

The performance of the cluster expansion model was optimized by minimizing the mean square deviation of the adsorption energies predicted by the model from the adsorption energies obtained by density functional theory (DFT) methods. In order to

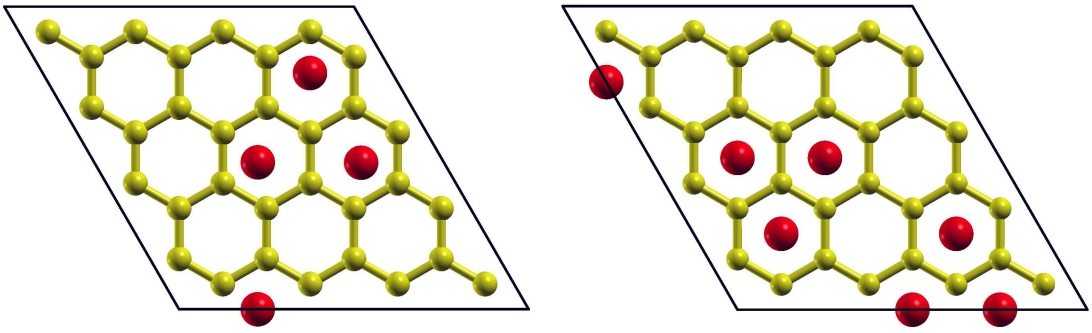

(a) (b)

Figure 3.8: A pair of generic Li atom configurations in a 4x4 graphene supercell with (a) 4 Li atoms per unit cell and (b) 7 Li atoms per unit cell.

avoid overfitting and favor locality, a penalty term was added to the cost function as described below. The total cost function per Li is defined as

$$\sigma_{Li}^2 = \frac{1}{N_{config}} \sum_{i=1}^{N_{config}} \left[\frac{E_{DFT}^i - E_{CE}^i}{N_{Li}} \right]^2 + t \sum_{i=1}^{N_J} J_i M_i J_i, \tag{3.7}$$

where the accuracy of the model is defined as

$$\sigma_{Li}^2 = \frac{1}{N_{config}} \sum_{i=1}^{N_{config}} \left[\frac{E_{DFT}^i - E_{CE}^i}{N_{Li}} \right]^2 = \frac{1}{N_{config}} \sum_{i=1}^{N_{config}} \sigma_i^2, \tag{3.8}$$

and σ_i^2 is the square deviation per Li of the model adsorption energy of configuration i from DFT one.

Similarly, the mean square deviation per carbon is given by

$$\sigma_C^2 = \frac{1}{N_{config}} \sum_{i=1}^{N_{config}} \left(\frac{N_{Li}}{N_C} \right)^2 \sigma_i^2. \tag{3.9}$$

In matrix form the cluster expansion for a given configuration C is given by

$$E_{CE}(C) = \sum_f J_f \Pi_f(C) \tag{3.10}$$

and the cost function per Li atom is given by

$$\sigma_{Li}^2 = \frac{1}{N_{config}} \sum_C (E_{DFT} - \Pi J)^T (E_{DFT} - \Pi J) + t J^T M J. \tag{3.11}$$

The penalty matrix M is a diagonal matrix which is defined so that the two body interactions are penalized proportionally to the distance of the interacting pairs; for instance, the element for second nearest neighbor interaction is multiplied by $(\sqrt{3.0})^\lambda$.

This imposed locality constraints to the 2-body terms similar to the one developed in Ref. [72]. The elements of M corresponding to the three and four body terms are set to 1.0. We used $t = 0.0005$ and $\lambda = 1.0$. Thus, the extra cost associated with the locality constraint is minimal and only play a significant role when less configurations than parameters are fitted.

The fitting was done by adding the configuration with the worst error to the training dataset starting from an initial configuration. In Fig. A.1 of Appendix A we show the evolution of the root-mean-square error as the number of training configuration increases. The complete discussion of the final model is discussed in Sec. 3.3.6.

3.3.4 Grand-canonical Monte Carlo

The finite temperature behaviour of Li interactions with graphene was investigated via Grand-Canonical Monte Carlo (GCMC) simulations which allow to compute the thermodynamic properties of the system at any given chemical potential, μ, and temperature, T. The GCMC results in this work were obtained with a 20x20 simulation cell of graphene containing 800 carbon atoms and 400 H-sites, while finite size effects were explored on a 25x25 cell.

The configuration space is sampled as follows: *i)* an adsorption site is selected at random; *ii)* an attempt to reverse its state is performed, i.e. adding a Li to the site if empty or removing it if occupied; *iii)* the proposed move is accepted according to Metropolis [70, 73] rule with probability $\min[1, \exp(-\beta(\Delta E - \mu\Delta N))]$, with $\beta = 1/(k_B T)$, $\Delta E = E_{new} - E_{old}$, $\Delta N = N_{new} - N_{old}$, and E_{new}, E_{old}, N_{new}, N_{old} are the configurational energies and number of adsorbed Li atoms of the new and the old configurations, respectively. Each attempt to add (delete) a Li to (from) a randomly selected site is considered a Monte Carlo step.

The temperature-chemical potential plane was explored with a series of unidimensional scans. Chemical potential scans were done between -0.2 and 0.2 eV at intervals of 10 meV while temperature scans were carried out between 400 and 4000 K at intervals of 50 K. In each scan, simulations were performed in sequence so that each simulation started from the final configuration of the preceding one. Thus, 1.5 million Monte Carlo equilibration steps were enough to equilibrate the system at each chemical potential/temperature pair. Statistics were then accumulated over 2 millions steps and used for the evaluation of the averages. When the simulation approached a phase transition, Li concentration, $x = \langle \sum_i c_i \rangle / N$, is observed to display sudden jumps and caused large uncertainties in the statistical averages calculated.

A few exploratory chemical potential scans at intermediate temperatures allowed us to identify the relevant phase transitions of the system and the corresponding chemical potential ranges by monitoring discontinuities in Li concentration. Then we performed

a careful chemical potential scan at very high temperature (4000 K) where all the phases are connected and no concentration discontinuity is observed. Subsequently, we perform temperature (annealing) scans at a number of fixed chemical potentials, which give us access to initial configurations for further chemical-potential scans for the different phases at various temperatures. In this way each (μ, T) state of a phase can be connected to any other (μ', T') state of another phase through a path in $\mu - T$ plane without discontinuous jumps of the relevant conjugate thermodynamical quantities, i.e. concentration and internal configurational energy. This allows us to compute the free energy of each phase on a common energy scale and compare their relative stability through thermodynamic integration.

3.3.5 Thermodynamic Integration

The total differential of the grand potential per site, $\Phi(\mu, \beta)$ can be used to obtain the following relation:

$$d(\beta\Phi) = \left(\frac{\langle E \rangle}{N} - \mu x \right) d\beta - \beta \, x \, d\mu. \tag{3.12}$$

where β is the inverse temperature $1/k_B T$.

At constant temperature, the integration of Eq. 3.12 from a reference chemical potential μ_0 to a desired chemical potential μ, gives the grand potential per site $\Phi(\mu, \beta)$ relative to the reference value, $\Phi(\mu_0, \beta)$, as:

$$\Phi(\mu, \beta) = \Phi(\mu_0, \beta) - \int_{\mu_0}^{\mu} x \, d\mu. \tag{3.13}$$

This expression is valid provided the chemical potential interval (μ_0,μ) spans a single phase and the resulting $\mu \to x$ mapping is invertible. The Helmholtz free energy per site of the phase under examination can then be obtained from the Legendre transform

$$F(x, \beta) = \min_{\mu} \{ \Phi(\mu, \beta) + \mu x \} = \Phi(\mu(x), \beta) + \mu(x) x. \tag{3.14}$$

In order to compare the Helmholtz free energy of two phases at a given temperature, the difference between the reference grand potential values of the two phases, $\Phi(\mu_0^{(1,2)}, \beta)$, must be determined. This can be achieved by thermodynamic integration on the path schematically illustrated in Fig. 3.9 that connects phase space point `a` to `d`. Point `a` corresponds to the reference point ($\mu_0^{(1)}$,β) of phase 1, and point `d` corresponds to the reference point ($\mu_0^{(2)}$,β) of phase 2. Integration of equation 3.12 over a smooth

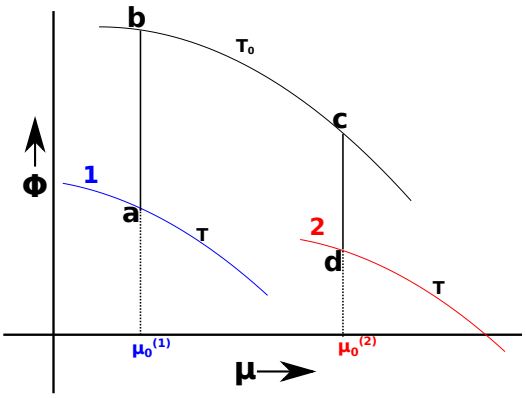

Figure 3.9: Thermodynamic integration: The free energy difference between phase 1 in a $(\mu_0^{(1)}, T)$ and phase 2 in d $(\mu_0^{(2)}, T)$ can be obtained using integration along the a $\rightarrow$ b $\rightarrow$ c $\rightarrow$ d path.

path without discontinuity is ensured by moving along the path a $\rightarrow$ b $\rightarrow$ c $\rightarrow$ d:

$$\begin{aligned}\beta\Delta\Phi^{(1,2)}(\beta) &= \int_{\beta}^{\beta_0}\left(\frac{\langle E^{(1)}\rangle}{N} - \mu_0^{(1)}\, x\right) d\beta \\ &- \beta_0 \int_{\mu_0^{(1)}}^{\mu_0^{(2)}} x\, d\mu \\ &+ \int_{\beta_0}^{\beta}\left(\frac{\langle E^{(2)}\rangle}{N} - \mu_0^{(2)}\, x\right) d\beta, \end{aligned} \tag{3.15}$$

where $\Delta\Phi^{(1,2)}(\beta) = \Phi(\mu_0^{(2)}, \beta) - \Phi(\mu_0^{(1)}, \beta)$ is the free energy difference at inverse temperature β between phase 1 at chemical potentials $\mu_0^{(1)}$ and phase 2 at chemical potentials $\mu_0^{(2)}$ and the auxiliary inverse temperature $\beta_0 = 1/(k_B T_0)$ is such that both the average energy, $\langle E\rangle$, and and the Li concentration, x, vary continuously along the path connecting the two phases. We found $T_0 = 4000\,\mathrm{K}$ to be sufficient for this purpose.

3.3.6 Li-graphene Interaction Potential

The values for parameters of the cluster expansion model are determined by fitting to the DFT adsorption energies (see Table 3.2). On the fitting set, a root-mean-square error of 18.0 meV/Li (1.6 meV/C) is obtained. The cross validation error calculated via leave-one-out method is found to be 30.2 meV/Li (3.4 meV/C). This error is significantly smaller than the root-mean-square deviation of 7 meV/C and cross validation score of 8 meV/C reported in Ref. [8]. The on-site ECI gives approximately the adsorption energy of an isolated Li ion, as expected. Among the 2-body interactions, the nearest neighbor 2-body ECI dominates, and is positive, in line with the strong electrostatic repulsion between Li ions in nearest hexagons, also in agreement with the cluster expansion parameters of Li-graphite system reported in Ref. [74]. The remaining two-body interaction terms, J_{2b} and J_{2c}, are negative, smaller in magnitude,

and decrease with increasing range. This supports the assumption that, once the long range electrostatic interaction is factored out, the cluster expansion can be truncated at few nearest-neighbor distances. Compact 3- and 4-body terms are found to be significant as well for a good description of the Li-graphene interaction. When neglected, a root-mean-square deviation of 98.3 meV/Li (7.5 meV/C) was obtained for the fit with a cross validation error of 158.2 meV/Li (21.2 meV/C). Similarly, when only the dipole-dipole long-range terms are neglected the model describes relatively well the Li-dense configurations, but dilute configurations with Li adatoms in 3x3 to 8x8 are poorly described indicating the importance of the long-range dipole-dipole interaction term. The root-mean-square deviation of 70.9 meV/Li (5.2 meV/C) with a cross validation error of 91.9 meV/Li (8.6 meV/C) was obtained. Plots and model parameters of these different scenarios are presented in Appendix A.3.

Table 3.2: Best fit values for the effective cluster interactions of the model in equation (3.5). All parameters, except γ, in eV

	J_1	J_{2a}	J_{2b}	J_{2c}	J_3	J_4	J_{dd}	γ
ECI	0.287	0.677	-0.109	-0.075	-1.429	0.234	2.086	0.383

3.3.7 Energetics of Lithium-Graphene Interactions

Following the procedure detailed in sections 3.3.5 and 3.3.4, we identify transition between two stable phases by observing discontinuities in the concentration as a function of μ at fixed temperature. In Fig. 3.10 we show an example of discontinuous transition between phases as observed at 800 K and at 4000 K where a continuous transition between phases can be observed in the GCMC simulations. Three notable phases are observed (see Fig. 3.11). These phases are referred in the following as the Li-gas (G) (Fig. 3.11a), the Li-islands (I) (Fig.3.11b) and the Li-stripes (S) (Fig. 3.11c) phases. The free energy per site of these phases at 800 K, as determined by thermodynamic integration, is shown in the upper panel of Figure 3.12. The minimum free energy for Li-gas phase is negative, lower than the bulk Li reference value, meaning that a small concentration of Li adatoms can stably reside on clean single layer graphene, contrary to the prediction of calculations at zero temperature. To examine the entropic origin of this stability, we compare the potential energy and entropy terms of the free energy of each phase at its minimum, i.e. at $\mu = 0$, as a function of temperature (see left panel of Figure 3.13).

At low temperatures, the energy term supersedes the entropic contribution in all the phases implying that free energy is positive, i.e. Li-graphene system is unstable with respect to bulk Li, in agreement with zero temperature calculations. At higher temperatures however, the entropic term dominates, stabilizing the Li-gas phase with

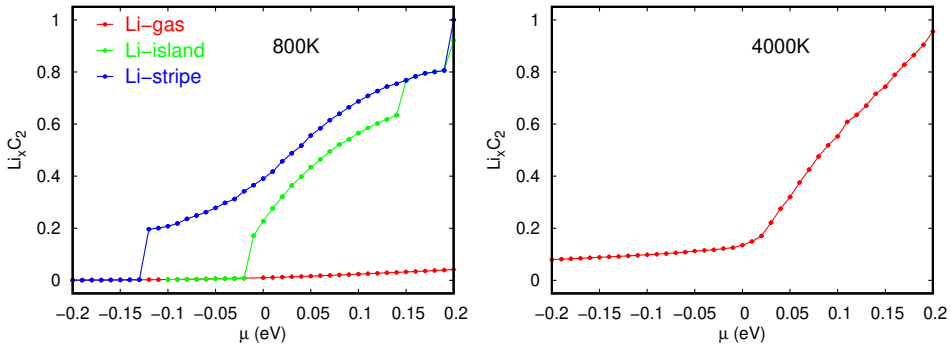

Figure 3.10: Concentration,x as a function of chemical potential, μ: The color codes indicate the phases. As we scan the chemical potential, we go from one phase (continues x) to another phase identified with a jump in the x. Here, In Fig. A.7 of Appendix A, we show the results from 800 K to 2000 K at a step of 200 K and include also the plot at 4000 K reference where a continues transition between phases can be observed.

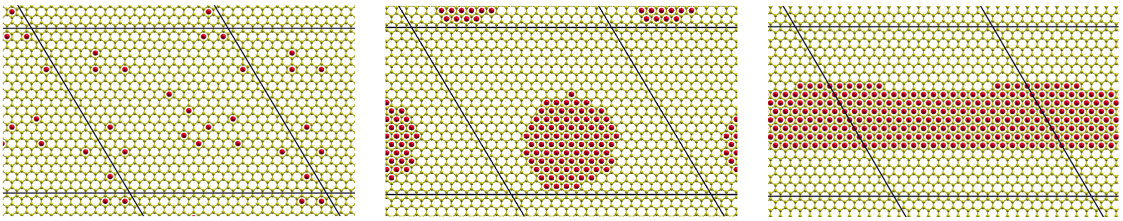

(a) Li-gas (b) Li-island (c) Li-stripe

Figure 3.11: A snapshot of Li ordering on graphene of the dispersed/gas phase, the Li islands, and the Li-stripes.

respect to bulk Li. In the Li-islands phase, there is a cross-over from an energy-dominated region (below 1200 K) to an entropy-dominated region (above 1200 K). In the stripe phase, energy is always the dominant term in the considered temperature range, free energy is always positive, therefore this phase is less stable with respect to bulk Li. We report in the right panel of Figure 3.13 that by allowing a modest 10 meV shift in chemical potential the energy-entropy balance in the various phases is significantly shifted to lower temperatures.

3.3.8 Li-graphene phase diagram

In order to identify the stable phases for a wide range of concentrations and temperatures, we construct the Li-graphene phase diagram using the common tangent construction. The common tangent construction at 800 K and the corresponding phase diagram boundaries are shown in the upper and lower panels of Figure 3.12, respectively. The maximum concentration at which the Li-gas phase is stable is about $x = 0.05$, equivalent to 1Li:40C at 2000 K, and it decreases progressively with temperature and becomes practically zero at 400 K. Below 400 K, the pure Li-gas phase becomes unstable and only a mixture of the Li-gas and Li-island exists. The region of coexistence between the Li-gas and the Li-island phases extends over a wide range, up to about $x = 0.3$

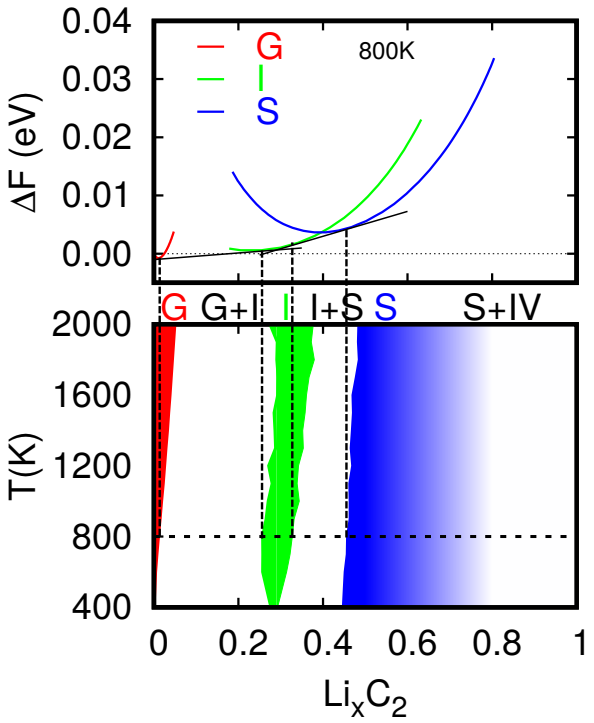

Figure 3.12: Free energy per site (upper panel) and the phase diagram of Li adsorption on graphene (lower panel). G, I, S and IV represent the gas(red), Li-islands(green), Li-stripes(blue) and islands-of-vacancies respectively. The region between the red and the green is the coexistence between the Li-gas and the Li-island phase while that between the green and the blue region is the region of coexistence between the Li-island and Li-stripes phases. The transition from the Li-stripes phase to island of vacancies as the concentration approaches to 1 is represented by the fading away of the blue region. The vertical dashed lines that extend across the panels are boundaries of the various coexistence regions at 800K.

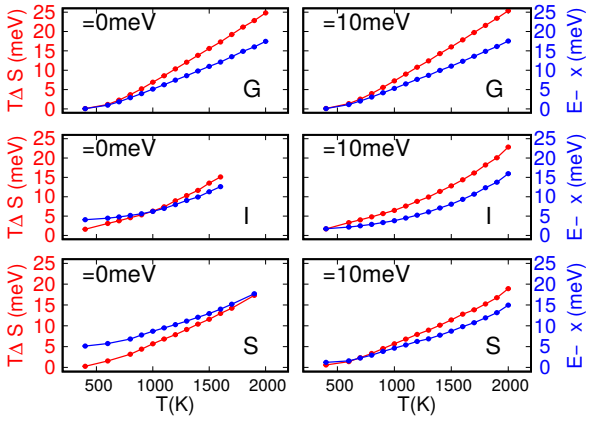

Figure 3.13: Entropy (red curve) and configurational energy (blue curve) contributions to the free energy as a function of temperature at chemical potential $\mu = 0$ and $\mu = 10\,\text{meV}$. $\mu = 0$ corresponds to the minimum of the Helmholtz free energy. G, I and S represent the Li-gas, the Li-island and the Li-stripe phase respectively.

(≈1Li:7C). Above $x = 0.3$, the Li-cluster phases are more stable at all temperatures. The stability of the Li-island phase occurs between Li:C≈1:7 and Li:C≈1:5. Between Li:C≈1:5 and Li:C≈1:4, the Li-island phase and the Li-stripe phase coexist in equilibrium. At about $x = 0.5$ (1Li:4C), the Li-stripe phase is the most stable.

The Li-stripe is observed to appear as a result of fusion of large Li-island clusters of adjacent simulation boxes. The stripe configuration hence reduces the surface of Li-clusters and provides an energetically more favorable growth path for increasing concentration, compared to continuous increase in the size of Li-islands. Obviously, in the thermodynamic limit, i.e. in a simulation box of infinite dimension, the fusion of two Li-islands would simply result in another Li-island of larger size and not in an infinitely extended stripe. Thus, in a bigger simulation box, the critical size of Li-island beyond which the Li-stripe will begin to develop is expected to be larger, while other phases such as Li-gas may remain largely unaffected by the size. In order to explore this finite size dependence, we performed simulations similar to the one described so far in a 25x25 graphene supercell, with 1250 carbon atoms and 625 adsorption sites per cell. The results are compared in Figure 3.14 where the stability domains determined in the larger cell are shown by hatched regions.

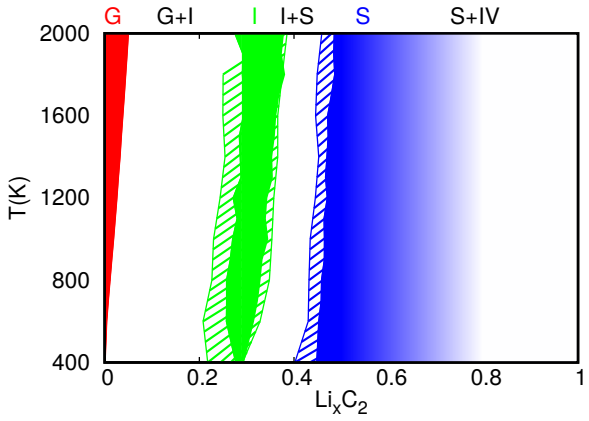

Figure 3.14: Finite size effect on temperature-concentration phase diagram. The solid shaded regions represent the phase diagram obtained in the 20x20 supercell while the hatched regions correspond to the one of the 25x25 supercell.

The Li-gas phase stability is found to be independent from the size of the simulation box implying that above 400 K and at very low concentration, Li ions will dispersed on graphene surface in the thermodynamic limit. We also observed an increase in the region of stability of Li-island phase with increasing cell size. Considering that in the bigger cell the same concentration is reached with more Li atoms (see Fig. 3.15), we see that in the bigger simulation box, bigger islands are stabilized with respect to stripe phase. This confirms the formation of Li-stripes as an artifact of the assumed periodic boundary conditions which would eventually disappear as the cell dimension becomes infinite.

The range of concentration of the phase coexistence regions are also observed to reduce with increasing cell size. This implies that the energy barrier between phases is reduced as the simulation box gets bigger and in the thermodynamic limit the 2D Li adsorption on graphene may occur continuously from Li-gas phase to fully lithiated

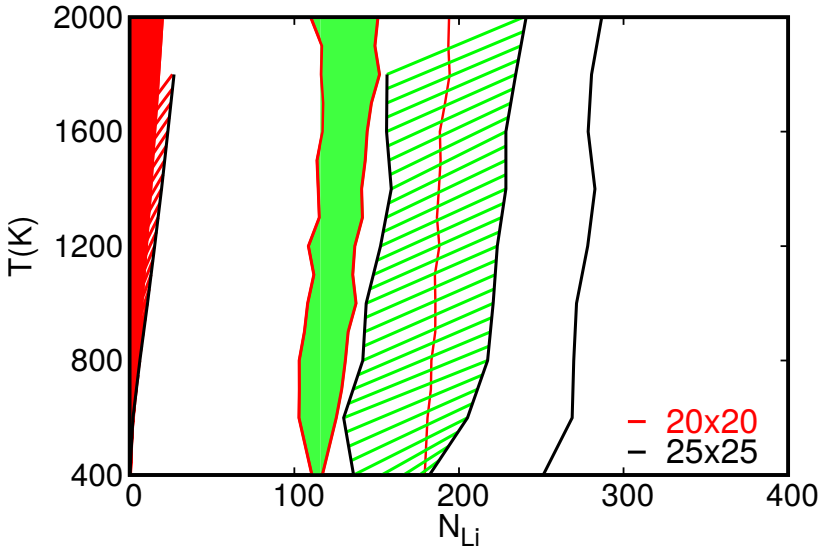

Figure 3.15: Finite size effects: Temperature versus number of Li, N_{Li} phase diagram. The red lines represent the phase boundaries determined with a 20x20 supercell and the black lines represent those determined with the 25x25 supercell. The red and green shaded region are the Li-gas and Li-island phases in a 20x20 cell while the red and green hatch regions corresponds to those obtained with 25x25 supercell. The plot shows the increased dimension of the Li-islands in equilibrium with the Li-stripe phase as a function of the similation box size.

graphene through the path of gas phase, gas-islands coexistence, Li-island and finally to complete lithiation respectively as the Li concentration increases.

Other phases such as islands-of-vacancies and gas-of-vacancies in an otherwise completely Li-covered graphene were also observed, however these phases are found to be very high in configurational energy. For instance, the energy required to create a single vacancy from a fully-lithiated graphene according to our model is 1.672 eV which is about six times the energy required to add a single Li atom to an empty graphene sheet, 0.287 eV. We can estimate the temperature needed to stabilize the gas-of-vacancy phase from the fact that the Li-gas phase is stable at 800 K and above. Assuming a simple proportionality gives a temperature of about 4800 K.

3.4 Conclusions

In summary, we have presented a study of the interaction of Li with single layer graphene at zero temperature by Density Functional Theory (DFT) and found that Li-graphene interaction is dominated by electrostatics due to charge transfer from Li to graphene. The trend of the adsorption energy per Li was found to be qualitatively different for Li adatoms compared to Li-clusters. This difference was found to be mainly due to a drop in the transferred charge in the case of clusters. Li adatoms and small clusters considered were found to be unstable against bulk Li metal at zero temperature. Finite temperature effects were studied by a combination of cluster expansion

techniques and Grand Canonical Monte Carlo methods. We found that thermal effects are sufficient to prevent phase separation of the Li-graphene system into its constituents at low concentrations but above the stoichiometry of $\approx LiC_6$, formation of Li-clusters is energetically favorable with respect to random distribution of Li atoms at all temperatures.

Our study focuses on 2D adsorption. Our systematic approach predicts two stable phases (Li-gas and Li-cluster) and the range of concentration (i.e. the chemical potential) at which the two phases are in equilibrium. This coexistence region, that could not be captured by previous approximate finite temperature treatments, covers about one-third of the entire concentration range at low temperatures. Another phase where Li islands join and form continuous stripes was also observed during Monte Carlo simulations however, by varying the size of the simulation cell, we show that it is due to finite size effects and will eventually vanish in the thermodynamic limit. Although a simulation artifact in the case of an infinite graphene sheet, such a phase can be speculated to occur at the edges of a graphene flake and seed the lithiation.

While 3D configurations could in principle be considered and were previously theoretically shown to be more stable with respect to random distribution of Li atoms both by zero- [14] and approximate finite-temperature calculations [15, 16], our 2D study may help interpret some interesting experimental observations such as the evidence of enhanced Li capacity in spatially constrained graphene materials where Li uptake positively correlates with the average interlayer distance. For example Ref. [3] reports that increasing the interlayer distance from 3.6 Å to 4.0 Å increases the capacity from 320 mAh/g, less than graphite, to 784 mAh/g, more than twice of graphite. In such densely packed, spatially constrained graphene materials, Li coverage model is more compatible with 2D rather than 3D clustering, and the increase in Li uptake with increasing interlayer spacing is compatible with a gradual transition from 2D to 3D cluster formation. Therefore, combined with these experimental observations in the literature, our results suggest that high Li uptake in dense graphene materials may be the result of significant 2D Li clustering. The computationally demanding study of vibrational effects, such as population of flexural phonons of graphene that may further stabilize Li adsorption, is left for future work.

Chapter 4

Machine Learning Based Interatomic Potentials

4.1 Introduction

In chapter 3, we developed and used cluster expansion based interatomic potentials and grand-canonical Monte Carlo to study Li interaction with graphene at finite temperature. A temperature-concentration phase diagram reveals remarkable results and provides great insight on the role of temperature on Li graphene interaction as the concentration of Li increases. In addition, we discussed the compatibility of the results with experiments.

Due to the limitation of the cluster expansion techniques such as the requirement of a well defined underlying lattice with unambiguously defined occupation sites, the cluster expansion method is not suitable to study the interaction of Li with general carbon host. For instance, to construct a cluster expansion based interatomic potential to describe Li interaction with different polymorphs of carbon such as graphene, graphite, amorphous, nanotube, nanoribbon and graphene nanoflakes, an independent potential must be developed for each of the phases assuming an underlying lattice can be unambiguously defined for each of the polymorphs. With introduction of defects to the carbon host, the problem becomes intractable with such site based method. These limitations necessitate a shift towards a methodology that permits the construction of an accurate and yet affordable interatomic potential to describe the interaction of Li with all carbon environments without the necessity to pre-define occupation sites. To this end, we delve into machine learning techniques which have been demonstrated to be capable of providing an accurate representation of the energy for a given atomic configuration at a much lower computational cost when compared to *ab initio* methods. In this chapter, we start with a brief introduction to machine learning, followed by a review of the artificial neural networks machine learning models for constructing interatomic

potentials. The chapter will be concluded with the discussion of the implementation of artificial neural network (ANN) in our recently developed machine learning package, "Properties from Artificial Neural Network Architectures (PANNA)" [75].

4.2 Machine Learning Interatomic Potentials

Machine learning (ML) is a discipline that focuses on the study of computer algorithms that improve automatically based on experience. ML algorithms construct a mathematical model that learn from data. These algorithms are now widely used in various applications such as computer vision, voice recognition, fingerprint identification, email filtering etc.

Nowadays, ML techniques are becoming popular in computational material science especially in the development of interatomic potentials suitable to study extended systems [76, 77]. This is due to their capability to approximate the high dimensional potential energy surface accurately. Basically, given atomic coordinates a ML model is constructed that maps the atomic coordinates to an energy which is subsequently fitted to *ab initio* data.

However, a direct mapping using the atomic coordinates as input is error-prone because ML algorithms have no knowledge of symmetry and therefore cannot understand for instance that a rigidly translated molecule should have the same energy as the untranslated one. Therefore, symmetry compliance representation of atomic system is paramount. Some examples of such representations in literature are vectors with respect to local frames [78], symmetry-based descriptors [79–81], graphs [82], matrices [83], list of bonds [84], chemical formulas [85] or molecular structures [86].

Once a representation is chosen, fast and accurate interatomic potential can be constructed using machine learning models such as neural networks [87–89] or Gaussian processes [77,90,91]. In this case, the total energy of a given atomic configuration is obtained from the sum of atomic energy contributions with each atomic energy depending on the representation of the chemical environment of that atom. This decomposition was proposed in Ref. [79] by Behler and Parrinello in 2007 where they used artificial neural networks combined with atom centered symmetry descriptors to approximate the potential energy surface of bulk silicon. Since then, other ML approaches in literature include Gaussian Approximation Potential (GAP) [80, 92] in which the atomic energy of a given environment is written as a linear combination of the Gaussian kernel between the test environment and all the environments in the training dataset, Spectral Neighbor Analysis Potentials(SNAPs) [93] which is a linearized version of the GAP method in the descriptors with a recent extension to quadratic dependence [94] and the Moment Tensor Potentials (MTPs) [95] where the interatomic potential is a linear combination of some polynomial basis functions of atomic environments which are

built from a set of moment polynomials etc. A recent work comparing these methods concludes that GAP methods have the highest accuracy, but also the highest computational cost, increasing with the size of the training dataset [96]. SNAP and MTP use lower cost regression strategies to correlate the local atomic environment with its contribution to the total energy.

In this work, we adopt the ANN approach coupled with the atom centered descriptors of Behler and Parrinello [79] and its modifications [75, 97]. The ANN functions are flexible and can be computed efficiently. They are continuous and relatively cheap compared to other approaches highlighted above [80, 92–95]. Its derivatives can also be computed efficiently.

The rest of the chapter is devoted to the presentation of our recently developed open-source PANNA package to build atom-centered descriptors, train and test ANN model. We also provide an example of the capability of the code by reproducing the results of Ref. [97].

4.3 Properties from Artificial Neural Network Architectures (PANNA) package

Here we present the philosophy and workflow of our recently developed ANN code, PANNA: "Properties from Artificial Neural Network Architectures", hosted on gitlab [98] and recently published in Ref. [75].

The PANNA code is an open-source, user-friendly, ML-based code for interatomic potentials generation written in python. It takes advantage of the optimization and hardware flexibility of Tensorflow (TF) [99] engine that allows ANN models to be developed, using multiple CPU/GPU/TPU, on a large training dataset. The code permits the construction of different network architectures for each atomic type and allows different trainability states of layers within a network (i.e the parameters of network layers can be frozen while the network is trained). Different activation functions are implemented such as Gaussian, hyperbolic tangent, ReLU, radial basis functions and linear functions. A variety of cost functions are also permitted. Finally, a LAMMPS interface is distributed along with the package to allow the application of the model constructed. Currently, the Behler and Parrinello descriptors are implemented as discussed in Sec.4.3.2. The general workflow of PANNA code is shown in Fig. 4.1.

4.3.1 Artificial Neural Network Model

Artificial neural network (ANN) is one approach to construct a machine learning model. It is made up of interconnected neurons (or nodes) that transmit signals among each other. The feed-forward neural networks (FFNNs), a class of ANN, consist of an

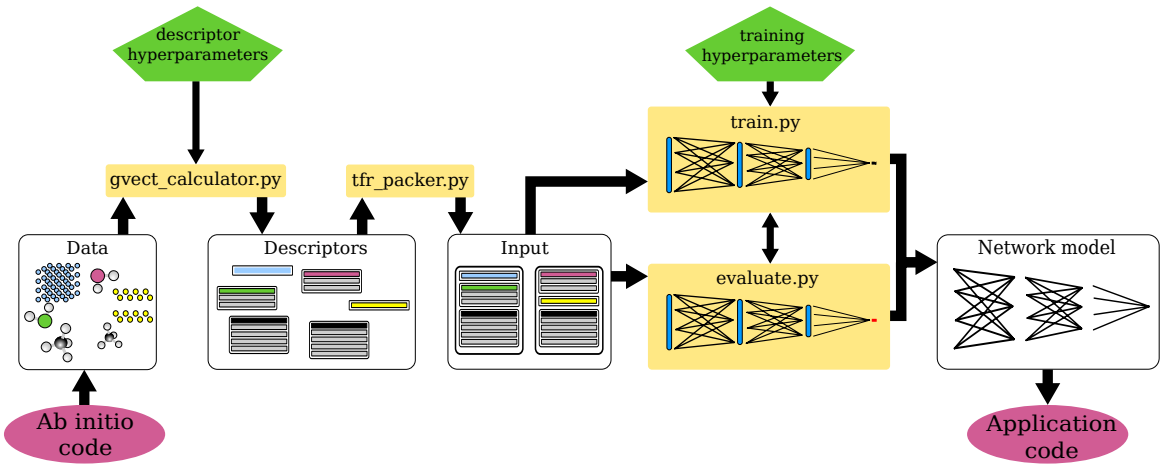

Figure 4.1: PANNA workflow [75]. Top layer (green) corresponds to user intervention in the neural network potential generation process. Input and network hyperparameters are determined at this layer. Currently several helper programs (not shown) included in the PANNA package aim to help users in this intervention. In principle, the hyperparameters are also part of the network optimization process and can potentially become an automated part of the workflow. Yellow boxes indicate main programs in PANNA while white boxes stand for the user-owned data in different stages of processing and the network potential. The interface with the third party codes are established via parser scripts and patches (not shown) included in the package.

input layer, one or more hidden layers (layers between input and output layers) and the output layer. The output layer is a single node that returns the prediction of an atomic energy in our application as shown in Fig. 4.2. Except for the input layer, the node, j, of a layer, $l+1$, receives signals from all the n_l nodes of the preceding layer, l, with a strength determined by the weight parameters, W^l_{ij}, connecting node j to i ($i = 1, ..., n_l$) which are then processed (a non-linear function of a linear combination of inputs) and passed as input to all the nodes of the next layer. This process is continued until signals are received by the output layer. Thus, the signal generated by a node j in layer $l+1$ is given by

$$a^{l+1}_j = g\left(\sum_{i=1}^{n_l} a^l_i W^{l+1}_{ij} + b^{l+1}_j\right) \tag{4.1}$$

where a^l_i is the i^{th} element of the output from layer l. The activation function for this node is specified by two parameters, the weight matrix W^{l+1} and the bias scalar b^{l+1}, and a non-linear function g.

This multilevel processing of the input layer with high connectivity among consecutive nodes of a fully connected FFNN yields a complex relationship between the input and the output node and therefore allows a high dimensional potential energy surface to be constructed.

Here, we adopt the Behler-Parrinello approach to atomistic neural networks [79] where the total energy of N-atom system is defined as the sum of atomic energy

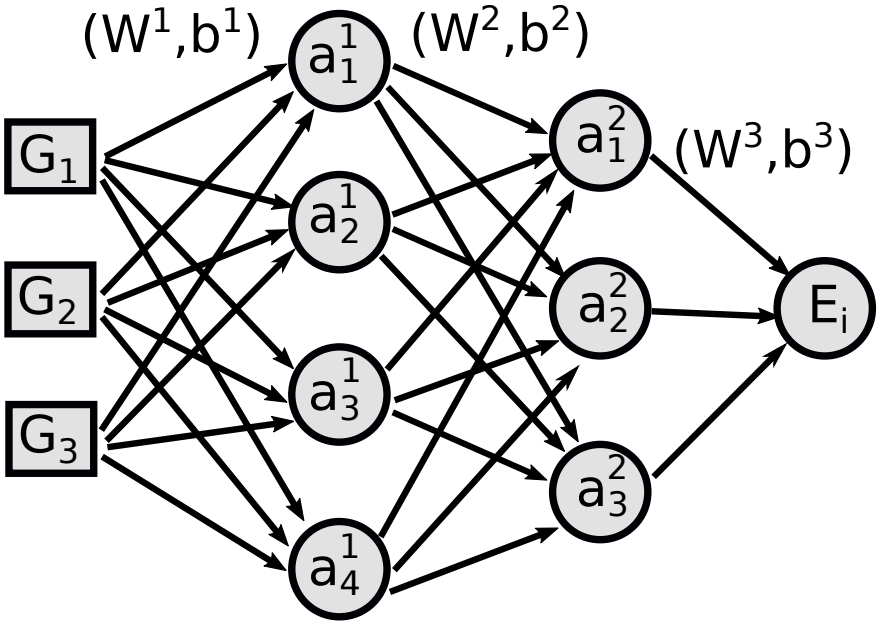

Figure 4.2: A schematic representation of an ANN with three inputs, two hidden layers and an output layer that give the atomic energy in this case. Each passing of information between layers l and $l+1$ is defined by the pair (W^{l+1}, b^{l+1}).

contributions:

$$E = \sum_{i=1}^{N} E_i(G_i), \tag{4.2}$$

where E_i is the energy contribution of an atom i, and G_i is its local environment descriptor vector discussed in Sec. 4.3.2.

To demonstrate how atomic energies are computed given a descriptor of an atom i, the number of hidden layers (and their width) and the non-linear activation functions, let us consider a feed-forward all-to-all connected ANN architecture depicted in Fig.4.2 in which the atomic environment of atom i is described by a vector of dimension 3 and fed to an NN architecture with two hidden layers of 4 and 3 nodes respectively, and an output layer with a single node. The network prediction of the atomic energy E_i reads as;

$$E_i = g_3(b_1^3 + \sum_{j=1}^{3} W_{1j}^3 \; g_2(b_j^2 + \sum_{k=1}^{4} W_{jk}^2 \; g_1(b_k^1 + \sum_{m=1}^{3} W_{km}^1 G_m))) \tag{4.3}$$

where g_1 ,g_2, g_3 are non-linear functions in general called the activation functions.

Once the ANN model is constructed, the weights and biases are determined from minimization of an objective function that measures the distance between the predicted value and the reference one provided the non-linear function g in equation 4.1 is differentiable with respect to its argument. In practice, a linear activation function is conventionally used for the output layer to allow arbitrary shift in the energy while Gaussian or hyperbolic tangent functions are commonly used for the hidden layers.

4.3.2 Atomic Environment Descriptors

Currently, two different types of descriptors are implemented in PANNA: atomic environment descriptors originally proposed by Behler and Parrinello in 2007 [79] and their modifications proposed in Ref. [75, 97]. Both descriptors represent the chemical environment of an atom by a fixed-size vector which depends on all other atoms within a user-defined spherical cutoff.

The functions include a two body and a three-body term, referred to as radial and angular descriptors respectively.

The radial descriptor function is defined as:

$$G_i^{\mathrm{Rad}}[s] = \sum_{j \neq i} e^{-\eta(R_{ij}-R_s)^2} f_c(R_{ij}), \tag{4.4}$$

where η and a set of Gaussian-centers R_s are user-defined parameters of the descriptor. The sum over j runs over all atoms whose distance R_{ij} from the central atom i is within the cutoff distance R_c. The cutoff function, f_c allows a smooth decay to zero at the cutoff radius and it is defined as:

$$f_c(R_{ij}) = \begin{cases} \frac{1}{2}\left[\cos\left(\frac{\pi R_{ij}}{R_c}\right)+1\right] & R_{ij} \leq R_c \\ 0 & R_{ij} > R_c. \end{cases} \tag{4.5}$$

The angular part of the descriptor with central atom i is defined as:

$$\begin{aligned} G_i^{\mathrm{Ang}}[s] =& 2^{1-\zeta} \sum_{j,k \neq i} (1+\cos(\theta_{ijk}-\theta_s))^{\zeta} \\ & \times e^{-\eta(R_{ij}/2+R_{ik}/2-R_s)^2} \\ & \times f_c(R_{ij}) f_c(R_{ik}). \end{aligned} \tag{4.6}$$

The sum runs over all pairs of neighbours of atom i, indexed as j and k, with distances R_{ij} and R_{ik} within the cutoff radius R_c, forming an angle θ_{ijk} with it. Here η, ζ, and the sets of θ_s and R_s are the user-defined parameters of the descriptor.

We note that the descriptor as written in Eq. 4.6 has discontinuous derivative with respect to atomic positions when atoms are collinear. To restore the continuity, we replace the $\cos(\theta_{ijk}-\theta_s)$ term with the following expression

$$2\frac{\cos(\theta_{ijk})\cos(\theta_s)+\sqrt{1-\cos(\theta_{ijk})^2+\epsilon}\sin(\theta_s)^2\sin(\theta_s)}{1+\sqrt{1+\epsilon\sin(\theta_s)^2}} \tag{4.7}$$

where we introduce a small normalization parameter, ϵ, such that the expression approaches $\cos(\theta_{ijk}-\theta_s)$ in the limit of $\epsilon \to 0$. In this work, $\epsilon = 0.001$ was used, while values between $0.001-0.01$ were found to yield stable dynamics and equivalent network

potentials for any practical purpose.

It is worth to mention that the sum over j in Eq. 4.4 runs through single species type at a time therefore making the radial part of the descriptors linear in the number of species. Likewise, the double sum in Eq.4.6 run over pairs of species type. The angular part of the descriptor scales qadratically with the number of species (i.e $n_s(n_s+1)/2$) for n_s species. The accuracy of a ANN interatomic potential depends strongly on the description of chemical environments. A good description of chemical environment is therefore needed to construct a flexible and transferable ANN interatomic potential.

In PANNA, the chemical environment descriptor are pre-computed on the whole training data before the commencement of training.

4.3.3 Network Architecture

In PANNA, a network is constructed for each atomic species such that for any atomic configurations atoms of the same type are evaluated with the same network which takes as input that atomic environment descriptors. In order to address the varying complexity in the environment of different species in an atomic system, PANNA allows each species network to be of different size (e.g different number of layers and different width of layers) . Furthermore, the option to freeze any layer of a species network is also implemented in PANNA. This is to enable a tighter control of the training dynamics and allows only a part of the network to be trained. Lastly, different non-linear activation function g can be chosen for each layer. Currently supported activation functions are:

- Gaussian: $g(x) = \exp(-x^2)$, as demonstrated in Ref. [97].
- ReLU (rectified linear unit): $g(x) = \max(0, x)$.
- Linear: $g(x) = x$
- Radial Basis Function (RBF) that changes the structure of the layer from the one outlined in Eq. 4.1 into $a_j^{l+1} = \exp\left[-\sum_{i=1}^{n_l}(a_i^l - W_{ij}^{l+1})^2\right]$.
- Hyperbolic tangent: $g(x) = \tanh(x)$.

4.3.4 Training

A ML model in general is trained by minimizing a loss function with respect to the parameters of the model (weights and biases). The loss function is a measure of the deviation of the model prediction from the reference value. The minimization of loss function in the PANNA package is done employing the commonly used Adaptive Moment Estimation (Adam) [100] algorithm based on stochastic gradient descent as implemented in tensorflow [99]. The loss function is computed on a randomly selected

subset (*minibatch*) of the training data at each optimization step which is the main source of stochasticity.

Let $W^{(t)}$ represent the parameters (weights and biases of all layers) of the ANN model at minimization time step t and let $O^c(W^{(t)})$ be the ANN prediction of an observable for a configuration, c, whose reference value is denoted by O^c_{ref}. Let $\mathcal{L}_c(W^{(t)}) = L(|O^c(W^{(t)}) - O^c_{ref}|)$ denotes the loss function on configuration c (a measure of the deviation of the ANN prediction for c from the reference). The loss function, $\mathcal{L}(W^{(t)})$, on a randomly selected batch at time step t is given by

$$\mathcal{L}(W^{(t)}) = \sum_{c \in \text{batch}^{(t)}} \mathcal{L}_c(W^{(t)}). \tag{4.8}$$

Let us denote $g^t_w = \nabla_w \mathcal{L}(W^{(t)})$ as the partial derivative of the loss function with respect to a given parameter w (weights or biases) and let $\mathbf{g}^t$ be a vector of the partial derivatives at time step t, then, the update of each parameter $w \in W^{(t)}$ at time step t in the Adam scheme depends on the previous gradients up to that at time step t. At each time step, t, an exponential moving average of the gradients, m^t, and its uncentered variance, v^t, read as follows

$$\begin{aligned} m^{(t)}_w &= \beta_1 m^{(t-1)}_w + (1-\beta_2) g^t_w, \\ v^{(t)}_w &= \beta_2 v^{(t-1)}_w + (1-\beta_2)(g^t_w)^2, \end{aligned} \tag{4.9}$$

where $\beta_1, \beta_2 \in [0, 1)$ are hyper-parameters that control the exponential decay rates of the moving averages. The smaller they are, the less history dependent the optimization. The moving averages are initialized to zero which biased the averages towards zero. In order to counteract this, the Adam scheme renormalizes the moving averages and equations 4.9 becomes

$$\begin{aligned} \hat{m}^{(t)}_w &= m^{(t)}_w / (1-\beta^t_1), \\ \hat{v}^{(t)}_w &= v^{(t)}_w / (1-\beta^t_2). \end{aligned} \tag{4.10}$$

The model parameter updates then reads as follows

$$w^{(t+1)} = w^{(t)} - \alpha \frac{(1-(\beta_2)^t)}{(1-(\beta_1)^t)} \frac{m^{(t)}_w}{\sqrt{v^{(t)}_w} + \epsilon'}. \tag{4.11}$$

where α is another hyper-parameter of the optimizer referred to as the learning rate and ϵ is a regularizer to prevent division by zero that may leads to stability issues.

The hyper-parameters β_1, β_2 and α are adjustable parameters. The best values are

those that allows the minimization of the loss function in few training time steps. In PANNA package, we allow an exponential decay of the learning rate, α, in the course of the training to mimic simulation annealing in the parameter space. The learning rate, α_t, at a time step t reads as follows;

$$\alpha^{(t)} = \alpha^{(t=0)} r^{t/\tau}, \tag{4.12}$$

where $\alpha^{(t=0)} = \alpha$ is the learning rate at step $t = 0$. The decay rate r and the decay step τ are used to determine the decay behavior which are chosen by the user.

In order to prevents overfitting, the L1- and L2-norm based regularizations of the model parameters are permitted in the PANNA package. The L1-norm imposes sparseness while the L2-norm prevents parameters from becoming excessively large. In PANNA, user-defined coefficients, γ_1 and γ_2 respectively for the L1- and L2-norm, are implemented that allow to scale the contribution of the regularization loss. The total regularization loss is written as

$$\mathcal{L}^R(W) = \gamma_1 |W| + \gamma_2 \frac{|W|^2}{2}. \tag{4.13}$$

The PANNA package allows the training of both energy and forces (the derivatives of the total energy with respect to atomic coordinates). While the training of the total energy of N atom systems imposes a single global constraints on the ANN model, force components provides $3N$ constraints. Therefore, it can be anticipated that training with force information will lead to a significant improvement in the accuracy of the model.

In the PANNA code, the prediction of atomic force on atom i, $\mathbf{F}_i$, is computed analytically via chain rule as

$$\mathbf{F}_i = -\sum_j \sum_\mu \frac{\partial E_j}{\partial G_{j\mu}} \frac{\partial G_{j\mu}}{\partial \mathbf{R}_i} \tag{4.14}$$

where $G_{j\mu}$ is the μ-th element of the descriptor for atom j. The derivative of the total energy with respect to the descriptors is already required during training for the update of network parameters (see equation 4.9) and the derivative of the descriptors with respect to atomic coordinates are pre-computed while the descriptors are being computed (see section 4.3.2).

The contribution of the force term to the loss function is implemented with a tunable parameter, c_F, that determines the relative importance given to the energy and forces during the optimization process. If we denote $\mathcal{L}_c^{i,\alpha,F}(W^{(t)})$ as the deviation of the prediction of force component α of atom i, from the reference one on configuration c

and $\mathcal{L}^F(W^{(t)})$ the loss on a minibatch, then we have

$$\mathcal{L}^{\mathcal{F}}(W^{(t)}) = c_F \sum_{c\in\text{batch}} \sum_{i=1}^{N_{\text{atoms}}} \sum_{\alpha\in\{x,y,z\}} \mathcal{L}_c^{i,\alpha,F}(W^{(t)}), \tag{4.15}$$

where c_F is a user-defined coefficient. If c_F is set to zero, only the energy is learnt.

4.3.5 Testing

After a model is successfully trained, one way to check the accuracy and transferability of the model is to evaluate it on independently generated configurations. In PANNA, the basic evaluation of the ANN model is implemented with the 'numpy' package as the evaluation of a network requires simple operations mainly matrix multiplication. In addition, models created with the PANNA package can readily be used, for example, to perform molecular dynamics in LAMMPS [101].

4.3.6 Programs

The construction and evaluation of ANN model with PANNA package is done with four main scripts. For each script, a configuration file based on 'python' configparser package which contains the variable of the script is taken as input. Here, we highlight the main features and capability of each script.

- `gvect_calculator.py`

  This python script is used to compute the atomic environment descriptors (and its derivatives with respect to atomic coordinate upon request) of a given configuration (referred to *example* or *simulation*) as discussed in Sec.4.3.2. The script stores the output in a single binary data file with the same file name as that of the example. The example must be written in PANNA-example format in standard JSON.

- `tfr_packer.py`

  This script is responsible for the collection and conversion of a large number of descriptor binary files into TFRecord (TFR) format which is ready to be efficiently processed by TF. The files produced by this script are referred to as TFData within PANNA. It reduces the I/O overhead, simplifies and speeds up the dataset management during training.

- `train.py`

  This script performs the task of training the ANN model. It reads the TFData files, handles the queue management to supply parallel processing of minibatches

and drives the training procedure with the appropriate TF calls. Among the parameters read by the script from the user is the maximum number of optimization steps train the model. Most often, there may be need to train the model for longer time step after the initial specified maximum steps is reached or due to an interruption of the training. This script stores the information required to restart the calculation as "checkpoints" during training in TF format at user-defined intervals. Additionally, the summary of each training is stored in TF "event" files that can be visualized in TensorBoard [99]

- `evaluate.py`

  The evaluation of the trained ANN model is done with this script. The script can parse the checkpoints, access the parameters of the network and evaluate the network prediction for a given input configuration at a user defined interval of the save checkpoints. It can operate on single binary descriptor files produced by `gvect_calculator.py` or bundled TFData files such as the outputs of `tfr_packer.py`.

For other technical details such as visualization of training using tensorboard, parallelization and a usage example, the reader is referred to Ref. [75].

4.4 Results: the case of molecules

In this section we present an example to demonstrate the usability of PANNA for aperiodic systems, with varying amount and quality of data. We report training of the network architecture previously used in Ref. [97]. The model contains a total of almost half a million parameters. This is a larger network than the average architecture employed in the literature. By reproducing the results of Ref. [97], we demonstrate that the training and testing modules of PANNA can answer the high performance demand scenarios in machine learning of interatomic potentials.

The dataset [102] used here contains 57462 small organic molecules consisting of H, C, N and O atoms, with up to 8 heavy atoms (atom species different from hydrogen) and corresponding total energy calculated via Density Functional Theory (DFT). The training dataset is partitioned into three subgroups, labeled as DSmax4, DSmax6 and DSmax8, each including data from molecules with up to 4, 6, or 8 heavy atoms, respectively.

The modified Behler-Parrinello symmetry functions described in Section 4.3.2 are used as in the original reference, with 32 Gaussian centers for the radial part, and 8 angular and 8 radial centers for the angular part. This choice of parameters results in descriptors of size $4 \times 32 + \frac{4\times5}{2} \times 8 \times 8 = 768$ for each atom. As it is used in Ref. [97],

the atomic network architecture consists of 3 hidden layers of sizes 128, 128 and 64, all with Gaussian activation function, followed by a linear activation layer to allow the arbitrariness in the energy. This gives rise to a model with 123265 parameters per species and a total of 493060 parameters. With almost half a million parameters, this is a larger network than the average architecture employed in the literature. By reproducing the results of Ref. [97], we demonstrate that the training and testing modules of PANNA can answer the high performance demand scenarios in ML of interatomic potentials. We note that the dataset contains no force information and therefore only the energy is trained upon.

Table 4.1 shows the size of datasets used for training and validation.

Table 4.1: The size of datasets obtained from Ref. [102] used for training and validation. As the dataset is constructed by sampling the normal modes of each molecule, alongside a scaling factor to reduce the bias towards bigger molecules, it contains a different amount of data for each molecule type, e.g. 480 examples for N_2 and 17280 for C_4H_{10} (butane) and 340 for C_8H_{18}(octane). The final models are benchmarked against 10347 configurations from normal mode sampling of 138 molecules from the GDB-11 [103] with 10 heavy atoms, also included in Ref. [102].

Label	Max # of heavy atoms	# of elements in training set [$\times 10^6$]	# of elements in validation set [$\times 10^6$]
DSmax4	4	0.656	0.134
DSmax6	6	3.432	0.427
DSmax8	8	17.476	2.182

In Ref. [97], a loss function proportional to the exponential of the square loss was proposed. In our test of this loss function, very large gradients were found in the early stages of the training leading to numerical instability. This instability was addressed by clipping the norm of the weight parameters. In this study a simple quadratic loss in combination with the capped exponential loss is found to alleviate the problem while still preserving increased gradients on the outliers. The total loss function reads as follows;

$$\mathcal{L}(W) = \sum_{c\in\text{batch}} \left(E_c^{\text{DFT}} - E_c(W)\right)^2 + \exp\left[a\tanh\left(\frac{1}{a}\sum_{c\in\text{batch}}\left(\frac{E_c^{\text{DFT}} - E_c(W)}{N_c}\right)^2\right)\right], \tag{4.16}$$

where E_c^{DFT} and $E_c(W)$ are the reference DFT energy and ANN energy prediction of configuration c with N_c number of atoms. $a = 5$ was used.

An initial learning rate of $\alpha = 0.001$ with a decay rate $r = 0.98$ and decay step $\tau = 3200$ is used following Eq. 4.12. In the case of training with DSmax8, the learning rate decay step is increased to $\tau = 16000$, leading to a slower decay. A fixed batch size of 1024 examples is used throughout.

In Figure 4.3 we depicts the evolution of root mean square error (RMSE) on training and validation datasets during training. For each dataset, three networks with identical architecture with different random initialization of weight and biases are trained. It can be seen that the proposed training schedule yields quantitatively reproducible results, which are consistently in good agreement with those reported in Ref. [97]. As it may be expected, the bigger the training dataset gets in variability, going from DSmax4 to DSmax8, the harder the training gets. Hence the training and validation errors increase steadily from approximately 0.3 to 1.2 kcal/mol per example, while the optimization steps required goes from approximately one to five million time steps.

The arithmetic average of different instances of trainings can be used to make committee predictions for each data point. The validation set RMSE resulting from such committees are 0.25, 0.52, 1.14 for DSmax4, DSmax6, DSmax8 respectively. Note that these values are very close to the best individual network prediction errors of 0.27, 0.57, 1.19 respectively (See also Fig. 4.3), indicating that energy prediction error of different networks for each example may be highly correlated. Single network predictions will be reported in the rest of this section.

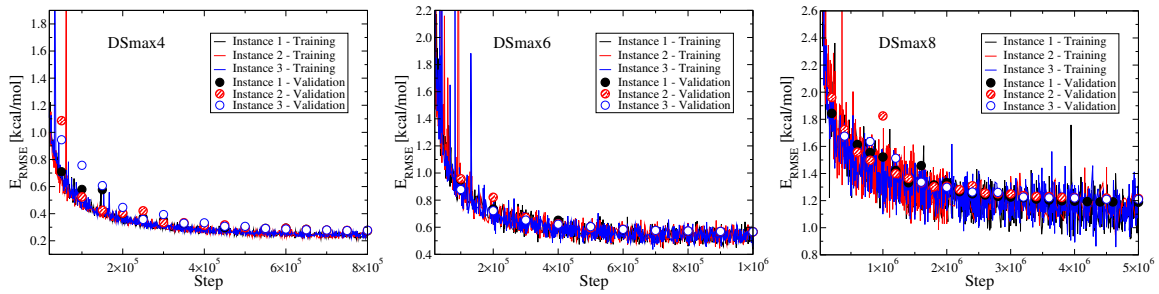

Figure 4.3: The energy RMSE during training and on validation set as a function of optimization steps for trainings with DSmax4 (left), DSmax6 (middle) and DSmax8 datasets (right). For each dataset, three instances of training is performed starting from different random initial parameters. The RMSE calculated at the final step for training (validation) set are 0.24 (0.27), 0.24 (0.28) and 0.24 (0.28) ; 0.54 (0.57), 0.54 (0.57) and 0.53 (0.57); and 1.12 (1.19), 1.18 (1.22) and 1.14 (1.21) in kcal/mol for trainings with DSmax4, DSmax6 and DSmax8 respectively. For comparison, the results from Ref. [97] are 1.16 (1.28) kcal/mol training (validation) RMSE for training with DSmax8.

While the analysis so far is based on training and validation sets drawn from the same distribution, the performance of neural network potentials can be better assessed based on their performance of configurations prepared from different distributions. To assess this property, we test networks on test sets of varying complexity. For example, a network trained on DSmax4 is tested on molecules with 5, 6, 7, 8 and 10 heavy atoms. In order to compare our results with Ref. [97] in the case of 10 heavy atom test set, the RMSE is also calculated with respect to the lowest energy structure for each molecule. RMSE calculated this way is referred as relative RMSE. The summary of the results is reported in Table 4.2. It can be seen from the energy-capped vs uncapped results that transferability is better in the low energy structures only compared to the high energy configurations. The overall network performance reduces as the training and test sets

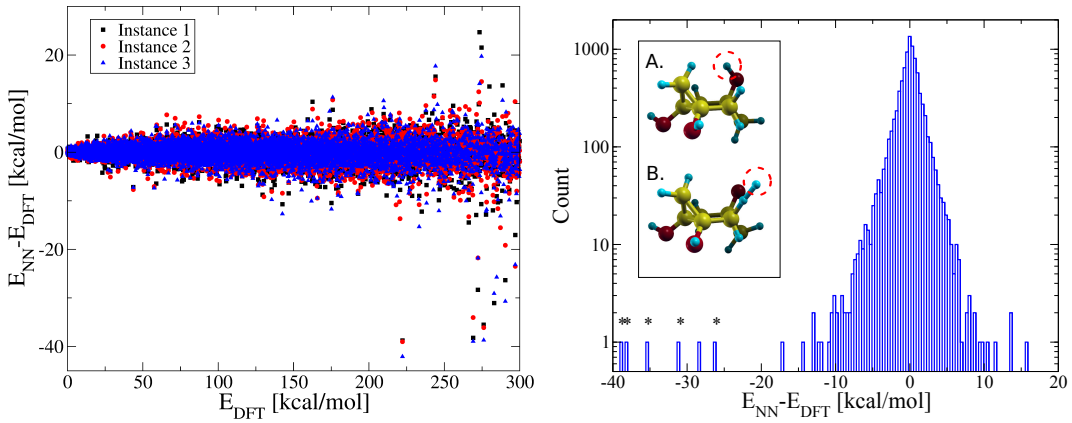

Figure 4.4: **Left:** Neural network prediction error compared to DFT energy on GDB-11 set with 10 heavy atoms where total energy is shifted so that the lowest energy structure of each molecule in the dataset corresponds to 0 kcal/mol. The networks are trained on DSmax8 set (see Fig. 4.3, bottom panel). **Right:** Prediction error distribution in log scale for Instance 2. It is noteworthy that all the marked outliers correspond to configurations of a single molecule shown in the inset where the H atom bound to O in the ground state configuration (A) is displaced far away, and may even form a bond with another H atom (B). Without the marked outliers, the RMSE reduces to 1.8 kcal/mol from 2.0 kcal/mol.

become more dissimilar.

Table 4.2: RMSE of energy prediction in kcal/mol for networks trained and tested with datasets of various molecular complexity. In the first five columns the RMSE is calculated for configurations where the total energy is within $E_{\text{cut}} = 275\,\text{kcal/mol}$ of the lowest energy configuration for each molecule. $DS\langle N\rangle$ stands for dataset with molecules having $\langle N\rangle$ heavy atoms. Additionally for the 10 heavy atom set, the RMSE for all configurations independent of their energies, and the relative RMSE for configurations within the lowest $E_{\text{cut}} = 300\,\text{kcal/mol}$ window is also given respectively. The last two columns contain the directly comparable values.

Training Set \ Test Set	DS5	DS6	DS7	DS8	DS10	DS10 all	DS10 300	Ref. [97]
DSmax4	17.1	23.2	28.3	30.0	24.5	139.2	21.0	26.0
DSmax6	0.5	0.7	13.5	15.5	14.5	138.5	15.5	17.7
DSmax8	0.6	0.7	1.2	1.4	2.1	87.7	2.0	1.8

Comparing the prediction of a network trained on DS8max with the DFT results per individual simulations shows that the error is larger for higher energy configurations (see Fig. 4.4). The distribution of error shows exponential decay for small error region with a visibly fat tail. Further investigations also show that the majority of the outlier configurations belong to a single molecule, hinting that careful error analysis beyond RMSE may be required for judging quality of network potentials.

Chapter 5

A Systematic Approach to Generating Accurate Neural Network Potentials: the Case of Carbon

This chapter is devoted to our recently proposed self-consistent approach that is based on crystal structure prediction formalism guided by unsupervised data analysis, to construct an accurate, inexpensive and transferable artificial neural network potential. We demonstrate the capability of this approach by using it to construct an interatomic potential for highly polymorphic elemental Carbon material and show its ability to reproduce first principles results on elastic and vibrational properties for diamond, graphite and graphene, as well as energy ordering and structural properties of a wide range of crystalline and amorphous phases.

5.1 Introduction

One of the most challenging and yet important step in constructing a general ML-based interatomic potentials is the training dataset generation. This become so because ML models learns the physics of the system through data presented to it during parameter optimization. Most often, training data are generated from known polymorphs of the systems of interest combined with molecular dynamics (MD) simulations to sample configurations in the basin of the chosen polymorphs [77, 91, 104, 105]. This approach strongly relies on a prior knowledge of all possible phases of the system which, in most cases, is not at our disposal. Elemental carbon is one example of a highly polymorphic material in which several different bonding types and structures exist. In this type of systems, construction of ML-based interatomic potentiasl from predefined polymorphs will certainly result in non-transferable potentials. This transferabilibity problem results in many different interaction models, each with limited applicability.

For example, among the several empirical force fields for carbon, the non-reactive, short range, bond-order-based Tersoff [106] model can describe dense sp^3 carbon structures while a highly parametric reactive force field (ReaxFF) [107] that explicitly includes long-range van der Waals interactions and Coulomb energy through charge equilibration scheme [108] is needed for structures with sp^2 hybridization. Furthermore, even though these empirical force fields give a qualitative understanding of materials properties, they are quantitatively inaccurate when compared to both *ab initio* methods and experiments [91, 104, 109, 110].

As already discussed in chapter 4, interatomic interaction models based on ANNs do not have a fixed functional form except for the network architecture (that has to be fixed), and their parameters are fitted to vast amounts of *ab initio* quantum mechanical data in the hope of assimilating the physics of the system into the parametrization. Hence the transferability *restraint* of classical force fields, that is due to their rigid form, is traded for a transferability *challenge* in the case of neural networks due to the (lack of) variety and completeness in the training dataset. To address this challenge of generating truly transferable ANN interatomic interaction models, training data must be obtained from an efficient and thorough sampling of the potential energy landscape. It should also be devoid of handcrafted initial configurations as much as possible. Such sampling of the very rugged and high dimensional landscape with *ab initio* electronic structure tools is a formidable challenge.

Here, we integrate evolutionary algorithm (EA) with molecular dynamics and clustering techniques in a self-consistent manner to sample the potential energy landscape and obtain data with high variability. The workflow we introduce extends the training data iteratively, similar to other active learning approaches that previously appeared in literature [76, 104, 111–113]. Unlike these methods that aim at constructing an optimal dataset for a specified part of the potential energy landscape, our workflow targets an unbiased training dataset which is necessary for increased transferability expected of a general purpose potential. Moreover, for reliable materials modelling, it is crucial to have indicators that signal when the limit of transferability is crossed. We address this aspect of ANN models by studying the relationship between data variability and transferability of the trained network via unsupervised data analysis. We demonstrate the performance of the approach highlighted above on the challenging example of crystalline and amorphous Carbon structures.

This study is a continuation of similar efforts in the literature: The first ANN interaction model for elemental carbon was developed in 2010 by Khaliullin *et al.* [104] to study graphite-diamond co-existence. The network was trained on an adaptive training set, where the starting configurations were manually selected from randomly distorted graphite and diamond phases, relaxed under a range of external pressures (from -10 to 200 GPa) at zero temperature. Then, new configurations obtained via

molecular dynamics simulation using the ANN potential constructed in the previous iteration step were added to the training data and a refined potential was constructed upon training on the new dataset. This procedure was continued until a self-consistency was reached in the prediction error on the new structures.

More recently in 2019, a hybrid model, where an ANN potential for the short-range interaction was supplemented with a theoretically motivated analytical term to model long-range dispersion, has been developed in order to address the properties of monolayer and multilayer graphene, with encouraging results [110]. As we will demonstrate in this work, ANN models such as these, built on data sampled solely from a limited part of the potential energy landscape can however be highly non-transferable. This transferability challenge for Carbon has been observed with kernel-based machine learning models as well.

Furthermore, a kernel-based model, specifically a Gaussian Approximation Potential (GAP), was constructed [91] using data from MD melt-quench trajectories of liquid and amorphous Carbon in 2017, to study amorphous structures. Motivated from its non-optimal behavior on crystalline phases, authors developed another GAP model with a specialized training data obtained via MD, for graphene [77]. Most recently, the same authors have developed yet another GAP interatomic potential for carbon with the aim to describe accurately all the known polymorphs of carbon with dataset obtained from different sources [114]. It is worthwhile to note that recently, a strategy combining kernel-based model generation with crystal structure prediction was suggested by Bernstein *et al.* [115].

However, since the computational cost for training and evaluation of a kernel-based model grows with the training set, this approach is suitable for small scale configuration space sampling only. In comparison, computational cost of neural networks are independent of the size of the training dataset, a feature that is exploited in the current study for accurate prediction of elastic and vibrational properties.

In this work we use a systematic approach to construct a highly flexible and transferable neural network potential (NNP) and demonstrate its application to the development of a general NNP for Carbon. We compare its performance with respect to other potential models previously optimized for specific phases and discuss the implications of our results for the trade-off between transferability and specialization.

5.2 Methods

5.2.1 Evolutionary Algorithm for Configuration Space Search

In iterative schemes, having a good starting point often means that a smaller number of iterations is needed to reach convergence. In a realistic use case scenario of NNPs, it

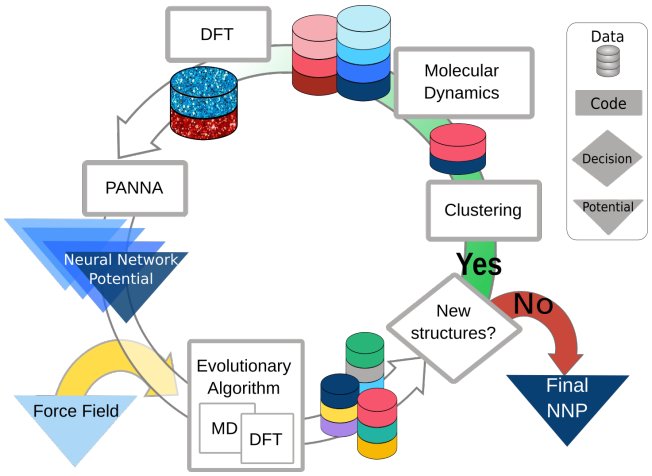

Figure 5.1: Sketch of the self-consistent scheme to generate an accurate and transferable neural network potential. The initial step to start the process (yellow arrow) can be performed with a classical force field as shown here, or any comprehensive dataset of structures such as the ones in Aflowlib [116], Materials Genome Initiative [117] or Nomad [118] repositories can be used to generate the first neural network potential model (blue triangle) to be refined through the self-consistent cycle. Once an initial potential model is chosen, Evolutionary Algorithm enables a diverse set of structures to be sampled. The following clustering-based pruning of structures further ensures that no single polymorph biases the dataset, i.e. at each step only novel structures (red and blue disks for the particular step highlighted above) are to be considered, further refined, and added to the dataset. The subsequent MD simulations sample the potential energy surface of each polymorph. Finally DFT calculations performed on a subset of MD-sampled structures are added to the *ab initio* dataset obtained thus far. The *ab initio* dataset augmented this way is then used to train the next neural network potential model (a darker blue triangle), starting the next cycle of the self-consistent scheme until no new structures are found by the Evolutionary Algorithm.

is reasonable to expect that only a moderately well-fitting potential would be available as a starting point. To demonstrate this, we start the self consistent cycle using a Li-C ReaxFF model to generate the initial configurations. This model is fit to DFT results with van der Waals correction and its details are set to describe well Li-C environments and defective graphite but not the wide range of solid C polymorphs considered in this work. We generate the initial configurations with 16 and 24 Carbon atoms per unit cell at 0, 10, 20, 30, 40 and 50 GPa via EA as implemented in USPEX [119, 120]. At each pressure, we start with a population of 30 (50) randomly generated structures for the 16 (24) atoms per unit cell, and evolve it through the following evolutionary operations with the given ratios: heredity (two parent structures are combined) 50%, mutation (a distortion matrix is applied to a structure) 25%, or by generating new random structures 25%.

At each generation, structures are optimized in five successive steps: (a) constant pressure and temperature molecular dynamics at 0.1 GPa and 50 K respectively for 0.3 ps with time step of 0.1 fs, (b) relaxation of cell parameters and internal coordinates until force components are less than 0.26 eV/Å, (c) constant pressure and temperature molecular dynamics at 0.1 GPa and 50 K respectively for 0.3 ps with time step of 0.1 fs, (d) relaxation of cell parameters and internal coordinates until force components are less

than 0.026 eV/Å, and (e) a final relaxation of cell parameters and internal coordinates until force components are less than 0.0026 eV/Å.

Only the 70% most energetically stable "parents" were allowed to participate in the process of creating the new generation. In the heredity step, only sufficiently distinct structures (whose cosine distance, as defined in the next section, is greater than a given threshold) are considered as parents. This threshold is fixed at 0.008 in the first iteration, as it is small enough to allow deformed structures from the same polymorph to be parents. In order to enhance the diversity of the structures in the subsequent iterations, the threshold is increased to 0.05 so that the parents can be expected to be from different polymorphs.

Each structure search is evolved up to a maximum of 50 generations at the first iterations and 30 in the subsequent ones. The configuration space search performed this way produces a wide range of sp^2, sp^3 and mixture of sp^2 and sp^3 structures, including defective layered structures.

5.2.2 Clustering

Initially, an unsupervised, bottom-up, distance-based hierarchical clustering approach with single linkage is used on all structures obtained with EA to identify the unique polymorphs. In the later iterations, clustering is applied only to those structures where NNP prediction differs from DFT ground truth energy by more than 5 meV/atom. That way, polymorphs that are already well described by NNP are not over-sampled. During clustering, to measure the similarity between structures, we use the fingerprint-based cosine distance defined in Ref.s [121] and [122]. In the case of a single species in the unit cell, and in its discretized form, the fingerprint of a configuration becomes:

$$F[k] = \frac{1}{2} \sum_{i,\text{cell}} \sum_{j} \frac{\text{erf}\left[\frac{(k+1)\Delta - R_{ij}}{\sqrt{2}\sigma}\right] - \text{erf}\left[\frac{k\Delta - R_{ij}}{\sqrt{2}\sigma}\right]}{4\pi R_{ij}^2 \frac{N^2}{V}\Delta} - 1 \tag{5.1}$$

where the first sum runs over all atoms i in the unit cell and the second sum runs over all atoms j within a spherical cutoff radius R_{max}, and R_{ij} is the distance between atoms i and j. The numerator describes the integral of a Gaussian density of width sigma over a bin of size Δ. N is the number of atoms in the unit cell and V is the unit cell volume.

The cosine distance between structures 1 and 2 is defined as

$$D_{\text{cosine}}(1,2) = \frac{1}{2}\left(1 - \frac{\mathbf{F_1} \cdot \mathbf{F_2}}{|\mathbf{F_1}||\mathbf{F_2}|}\right). \tag{5.2}$$

The dimension of the F-vector is set to $R_{\text{max}}/\Delta = 125$ with $R_{\text{max}} = 10$Å and $\Delta = 0.08$ in this work. Two configurations closer to one another than a distance threshold are

determined to belong to the same cluster. In this work the threshold is tuned to yield approximately 100-150 clusters at each step, which results in affordable computational cost for the remaining calculations of the self-consistent cycle.

5.2.3 Molecular Dynamics

We manually select a representative structure from each cluster and perform a 0.5 ns classical NPT molecular dynamics simulation with Nose-Hoover thermostat and barostat. In these simulations the external conditions of pressure and temperature are ramped up from -50GPa at 100 K, to 50GPa at 1000K in the course of 0.5 ns. The characteristic relaxation times of the thermostat and barostat are chosen as 50 fs and 100 fs, respectively. By sampling a snapshot of the dynamics every 5 ps, 100 configurations are selected. All molecular dynamics simulations are performed with LAMMPS package [123]. In addition, 440 randomly selected graphene atomic configurations from the libAtoms repository [124] are added to the selection. This set constitutes the set of structures where *ab initio* total energy calculations are then performed and added to the training set.

5.2.4 First Principles Calculations

The first principles calculations performed on all the structures visited during EA configuration space search and MD refinement described earlier employ the following parameters: Plane wave basis set kinetic energy cutoff for wavefunctions and charge density are 80 and 480 Ry respectively. The rVV10 [26] exchange-correlation functional that incorporates non-local van der Waals correlations is employed. A Brillouin zone sampling with resolution of $0.034\times2\pi$ Å^{-1} for the 3D carbon structures and $0.014\times2\pi$ Å^{-1} for graphene is used. These parameters are found to yield 1mRy/atom precision on diamond, graphite and graphene. All DFT calculations were performed with the Quantum ESPRESSO package [51, 52]. Elastic properties are computed through the thermopw framework [125] while vibrational properties are obtained using PHON package [126] with forces computed with LAMMPS package [123].

In the first self-consistent iteration, the training set is made up of all generated structures lying within 10 eV from the lowest energy one. This results in a total of ∼16000 configurations. In the subsequent iterations of the self-consistent procedure, we use all configurations whose energy per atom is within 1.2 eV of the lowest one, these are added to the previously selected configurations, amounting to a total of about 30000 configurations in the second and 60000 configurations in the third and final iteration. From these configurations, 20 % was set aside for validation and the remaining 80 % was used in the NNP training.

5.2.5 Neural Network Architecture

The artificial neural network (ANN) described in chapter 4 is used as implemented in the PANNA package also summarized in chapter 4. We choose descriptors of dimension 144 per atomic environment. The contribution of an atom to the total energy is obtained by feeding its environment descriptor to the feed-forward all-to-all-connected neural network. Here we build a network with two hidden layers, with 64 and 32 nodes for the first and second layer respectively, both with Gaussian activation function, and a single-node output layer with linear activation. The resulting network has a total of 11393 parameters, i.e. $(144 \times 64) + (64 \times 32) + (32 \times 1) = 11296$ weights and $64 + 32 + 1 = 97$ biases. The energy of each atom is then summed to obtain the total energy of the configuration. The force on each atom can be obtained analytically following the description in Sec.4.3.4 of chapter 4.

During training, the weight and bias parameters W, are optimized with the Adam algorithm [100] using gradients obtained by randomly selected subsets (minibatches) of data. The loss function of this stochastic optimization problem is defined as the sum of two contributions: one using the total energy value (Eq. 5.3) and one using the force on each atom (Eq. 5.4):

$$\begin{aligned}\mathcal{L}^{\mathrm{E}}(W) = &\sum_{c\in\mathrm{batch}} \left(E_c^{\mathrm{DFT}} - E_c(W)\right)^2 \\ &+ \exp\left[a \tanh\left(\frac{1}{a}\sum_{c\in\mathrm{batch}}\left(\frac{E_c^{\mathrm{DFT}} - E_c(W)}{N_c}\right)^2\right)\right],\end{aligned} \tag{5.3}$$

where E_c^{DFT} is the ground-truth total energy obtained via DFT and E_c is the NN prediction for total energy of a given configuration c, consisting of N_c atoms in the unit cell. The second part of this equation exponentially penalizes outliers while keeping the exponent normalized; a is a constant that allows to tune this penalty; $a = 5$ is used in this study. The force contribution to the loss reads as follows:

$$\mathcal{L}^{F}(W) = \gamma_F \sum_{c\in\mathrm{batch}} \sum_{i=1}^{N_c} \left|\mathbf{F}_i^{\mathrm{DFT}} - \mathbf{F}_i\right|^2, \tag{5.4}$$

where for any atom i of configuration c, $\mathbf{F}_i^{\mathrm{DFT}}$ is the ground-truth force obtained via DFT, and $\mathbf{F}_i$ is the NN prediction for it. γ_F is a user-defined parameter that controls the scale of this loss component. The results reported are obtained with γ_F equals 0.5. The relative error loss highlighted in Results section is defined as

$$\mathcal{L}^{F}(W) = \gamma_F \sum_{c\in\mathrm{batch}} \frac{1}{N_c} \sum_{i=1}^{N_c} \frac{\left|\mathbf{F}_i^{\mathrm{DFT}} - \mathbf{F}_i\right|^2}{\left|\mathbf{F}_i^{\mathrm{DFT}}\right|^2 + f_0^2}, \tag{5.5}$$

where f_0 is a regularizer constant, chosen as $f_0 = 260\,\mathrm{meV/Å}$ in this work.

An L_2-norm regularization term is also added with a small coefficient $\gamma_R = 10^{-4}$ to prevent weights from becoming spuriously large:

$$\mathcal{L}^R(W) = \gamma_R \frac{|W|^2}{2}. \tag{5.6}$$

The total loss is thus defined as

$$\mathcal{L}(W) = \mathcal{L}^{\mathrm{E}}(W) + \mathcal{L}^F(W) + \mathcal{L}^R(W). \tag{5.7}$$

All models are trained starting from random weights and a starting learning rate $\alpha_0 = 0.001$. The learning rate is decreased exponentially with optimization step t following equation 4.12 in Sec.4.3.4 of chapter 4 with decay rate $r = 0.96$ and the decay step $\tau = 3200$. A batch size of 128 data points is used throughout the study.

To ensure stability of the angular part of the atomic environment descriptors, a regularization term was introduced (see Sec.4.3.2 in chapter 4). Based on experience, values of ϵ parameter between $0.001 - 0.01$ yield stable dynamics. Here, we used $\epsilon = 0.001$.

The radial descriptors are parametrized with $\eta = 16.0$ Å^{-2}, while 32 equidistant Gaussian centers, R_s, are distributed between 0.5 Å and 4.6 Å. For the angular part $\eta = 10.0$ Å^{-2}, $\zeta = 23.0$, 8 equidistant R_s are distributed between 0.5 Å and 4.0 Å and 14 θ_s are chosen between $\pi/28$ and $27\pi/28$ with spacing $\pi/14$. The cutoff R_c is 4.6 Å for radial and 4.0 Å for the angular descriptors, respectively. The resulting descriptor has a total of $32 + 14 \times 8 = 144$ components per atomic environment.

5.3 Results

5.3.1 Self-Consistent Training and Validation

The neural network potential(NNP) is constructed following the self-consistent approach sketched in Fig. 5.1. This recursive data-creating and fitting cycle starts with a trial force field (FF) which is used to generate an initial set of configurations via evolutionary algorithm. In the absence of an established FF model for a new material, rough approximations such as Lennard-Jones or low-cost DFT approximations can be used with small unit cells for the very first iteration. Evolutionary algorithms (EA) are commonly used in crystal structure prediction studies as they allow efficient sampling of the configuration space. Their success in thorough sampling is demonstrated by their ability to predict new crystal structures before the experimental observation [127, 128]. As the exploration of the configuration space continues, a single point DFT calculation is performed on each distinct polymorph generated by EA. These structures are then clustered using a distance measure. From each cluster, a representative example is manually selected and a classical molecular dynamics simulation at a given pressure and temperature range is performed. The additional MD simulation step allows the sampling of the whole neighborhood of the equilibrium configuration for each polymorph, resulting in accurate prediction of structural properties for every polymorph. The dataset obtained this way is used to train a neural network model. The trained NNP is then used for starting a new iteration of the self-consistent cycle. This increases the training set diversity, by preventing the energetically favorable structures that are easily accessed by EA from dominating the whole training set. The iterative procedure highlighted above is repeated until no new structures are found.

While iterative expansion of training set is not a new idea, our implementation pushes its limits in diversity and balance: we use a full EA to sample configurations, without anchoring the search in any known polymorph or rigid transformations between polymorphs as in Refs. [112] or [113]. This makes our method applicable to novel materials with unexplored phase space and prevents any bias towards known phases. We then use clustering, which allows to achieve a balanced set despite the tendency of EA to sample stable configurations more often. Finally, starting from a representative configuration for each cluster, we perform MD simulations so that equilibrium properties of every polymorph are well described independent of their stability with respect to the ground state. We refrain from using active learning methods that depend on network agreement (as in Ref. [76]) as network prediction errors are not guaranteed to be uncorrelated, e.g. two networks may agree on the wrong result, especially if under-parameterized. We also refrain from expanding the training set with structures obtained solely through MD trajectories as in Ref. [129], because of the risk of missing significant polymorphs that would only be sampled rarely, and with

decreasing frequency, i.e. requiring longer and longer MD runs to run into significant additions to the dataset. Instead, a coherent integration of EA, clustering and MD together yields an unbiased, balanced and diverse dataset.

Table 5.1: Training and Validation RMSE: The training RMSE is the average over the batch RMSE of the last 2500 training steps, while the validation RMSE is evaluated over the entire validation set with the NNP obtained at the last training step. The energy (force) mean absolute errors on the validation set are 4 (0.09), 12(0.12) and 14(0.16) meV/atom (eV/Å) for first, second and third iteration respectively.

No. of iterations	Size of data	Energy RMSE (meV/atom)		Force RMSE (eV/Å)	
		train	validation	train	validation
1	15841	6.8	6.8	0.14	0.14
2	30815	17.1	20.0	0.19	0.22
3	60133	22.0	22.1	0.26	0.27

The performance of an NNP at each self-consistent loop is evaluated during training via the validation scheme. Fig. 5.2 shows the evolution of NNP energy accuracy on the training and validation set as a function of training steps at each self-consistent iteration (first row). The training root-mean-square error (RMSE) corresponds to the instantaneous RMSE computed on the elements of the batch considered at that training step while the validation RMSE is computed on all the configurations in the validation set. The RMSE on the validation set agrees with the training RMSE throughout the training, an indication that the model does not overfit to the training data set. The analysis of the force prediction error at different stages of training gives similar results and are depicted in Fig. 5.3.

The increase in energy and force RMSE from iteration 1 to 3 is a result of the increase in the diversity of the atomic environments. At each self-consistent iteration the diversity of the dataset increases as new structures are explored (see Table 5.1), while the number of parameters of the network, therefore its flexibility, is kept fixed. It is worthwhile to note that the prediction error for both energy and forces is not distributed according to a Gaussian distribution function but a fatter-tailed one (see second row of Fig. 5.2 and 5.3). Therefore, while the RMSE given here is a good measure to compare training and validation error with one another, it overestimates the average NNP prediction error in general.

To demonstrate how the general accuracy of the NNPs is changing with each iteration, we check their performance on a dataset of 197 distinct carbon structures. These structures were obtained by Deringer and co-workers [130] via random search of crystal structure of carbon with a GAP developed for liquid and amorphous carbon systems [91] and are distributed online [124]. For consistency, their energies are recalculated with the same DFT parameters as explained in Methods section. Fig. 5.4 shows the energy ranking as predicted by NNP, GAP, Tersoff and ReaxFF. It can be

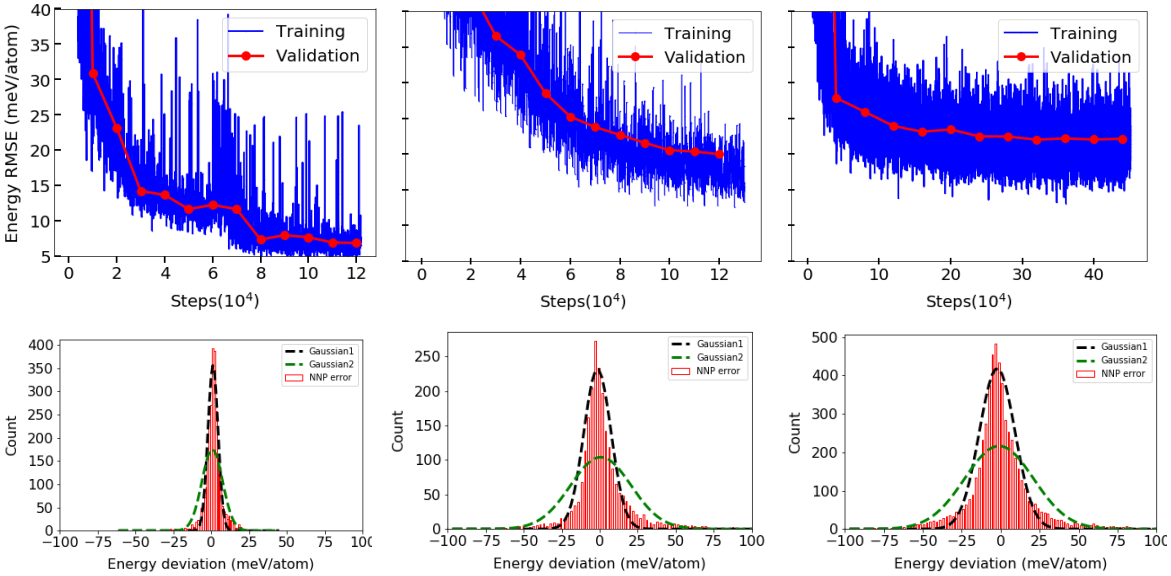

Figure 5.2: Training and validation error in energy. **First row**: The evolution of error in energy on training and validation set for potentials trained at first (*left*), second (*middle*) and third iteration (*right*) of the self-consistent cycle. The blue lines are the RMSE on a given batch of 128 configurations at the corresponding time step during training. The networks are evaluated during training on all the validation set of sizes ≈3000, ≈5200 and ≈12000 configurations for first, second and third iterations, respectively (red dots with lines as guide to eye). **Second row**: Error distribution for the validation dataset at first(*left*), second (*middle*) and third (*right*) iteration respectively. The histogram is constructed such that the area is equal to the number of configurations. The black dashed line is a normalized Gaussian fit, resulting in a mean $\mu = 1.19$ meV, $\mu = -1.43$ meV, $\mu = -2.22$ meV and standard deviation per atom of $\sigma = 3.36$ meV, $\sigma = 8.94$ meV, $\sigma = 11.44$ meV, respectively clearly failing to fit the fat tailed distribution. The green dashed line are the Gaussian distributions obtained with the mean ($\mu = 0.72$ meV, $\mu = 0.74$ meV, $\mu = -1.02$ meV) and standard deviation per atom ($\sigma = 6.85$ meV, $\sigma = 19.92$ meV, $\sigma = 22.13$ meV) of the error distribution in the validation set corresponding to distributions with wider spread.

seen that the NNP accuracy gets better with each iteration. The third iteration NNP accuracy agrees remarkably well with DFT results and performs better than all the other methods tested. It is noteworthy that the final NNP carries no signature of the ReaxFF used in the initial step to explore the configuration space. Both classical potentials, Tersoff and ReaxFF, perform very poorly compared to machine learnt ones, and the NNP outperforms GAP results published in Ref. [91,130], albeit GAP was fitted on *ab initio* data obtained with LDA exchange correlation functional [132]. For fair comparison, we train a new NNP, using the same training dataset obtained via the self consistent procedure, but with LDA functional. This potential, referred as NNP-LDA, performs similarly to the NNP highlighted in this work, and similarly outperforms all the other potentials. In the rest of the work, the results denoted with NNP refers to the potential that is trained with the rVV10 functional unless otherwise specified.

Table 5.2: Elastic properties of Diamond, Graphite and Graphene. For Diamond (**top**) all machine learnt potentials reproduce their reference DFT lattice parameter with less than 1% relative error. When the comparison extends to elastic constants and bulk modulus, however, only the NNP described in this work shows a consistently close agreement between DFT and the potential model, $< 3\%$ relative error, the range of variation also observed between two different experiments. For Graphite (**middle**) we report the lower bound for the bulk modulus using Reuss average, i.e. $1/B_0 \equiv s_{11} + s_{22} + s_{33} + 2(s_{12}+s_{23}+s_{31})$. The robust intraplanar structural features of graphite is captured well by all machine learnt potentials while the weaker interplanar interaction and, in particular, elastic properties that couple the two, are more challenging to capture. This is true even for the hybrid potential hNN-Gr of Ref. [110] where the distance dependence of the long range interaction is manually set to r^{-6} and the potential is tailor-fit to describe multi-graphene systems. For Graphene (**bottom**) the 2D elastic constants were computed with the normalized 3D stress as $\sigma_{2D}(\epsilon) = E(\epsilon)/A_0$ at a given strain of ϵ, where $E(\epsilon)$ is the total energy at ϵ and $A_0 = \sqrt{3}a^2/2$ is the area of graphene plane. E is the Young modulus and ν is Poisson's ratio. The elastic constant is converted to bulk properties in GPa by dividing by the interlayer distance $c/2$ of graphite reported in Table 5.2 or in respective experimental reference. Both machine learnt potentials similarly overestimate the Poisson's ratio and underestimate the Young's modulus with respect to their DFT references. Differences between DFT references are of similar magnitude as the differences between NNP and DFT in each case

	This work		Ref [91]		Ref [104]		Ref [106]		
Diamond	NNP	DFT	GAP	DFT	NNP	DFT	Tersoff	Exp [133]	Exp [134]
a (Å)	3.576	3.584	3.539	3.532	3.569	3.570	3.566	3.567	
B_0 (GPa)	431	425	438	466	434	439	426	442	445
C_{11} (GPa)	1054	1044	1090	1101	1016	1056	1074	1079(5)	1080
C_{12} (GPa)	119	116	112	148	142	130	102	124(5)	127
C_{44} (GPa)	542	547	594	592	580	567	641	578(2)	576

	This work		Ref [104]		Ref [110]		
Graphite	NNP	DFT	NNP	DFT	hNN-Gr$_x$	DFT	Experiment
a (Å)	2.471	2.471	2.467	2.467	2.467	2.466	2.464[b], 2.463[c]
c (Å)	6.732	6.719	6.688	6.815	6.804	6.800	6.712[b], 6.712[c]
B_0 (GPa)	48	40	48	37	-	-	36(11)[b]
C_{11} (GPa)	1053	1048	1080	1069	978	1080	1060(20)[a], 1109(16)[b]
C_{12} (GPa)	197	182	179	162	177	162	180(20)[a], 139(36)[b]
C_{13} (GPa)	-23	-5	0	-4	-67	-5	15(5)[a], 0(3)[b]
C_{33} (GPa)	57	43	52	40	40	33	37(10)[a], 39(7)[b]
C_{44} (GPa)	-5	4	7	5	1.79	3.36	0.27[a], 5(3)[b]
C_{66} (GPa)	428	433	-	-	-	-	485(11)[b]

	This work		Ref [110]		Ref [106]	
Graphene	NNP	DFT	hNN-Gr$_x$	DFT	Tersoff	Experiment
a (Å)	2.470	2.470	2.467	2.466	2.530	2.46[d]
ν	0.244	0.173	0.197	0.149	-0.158	-
E (GPa)	967	1015	1021	1060	1216	1015(149)[e], 2400(400)[f]
C_{11}(GPa)	1028	1047	1062	1084	1247	-
C_{12} (GPa)	251	181	209	161	-197	-

Experiments: [a]Ref [135] [b]Ref [136] [c]Ref [137] [d]Ref [138] [e]Ref [139] [f]Ref [140]

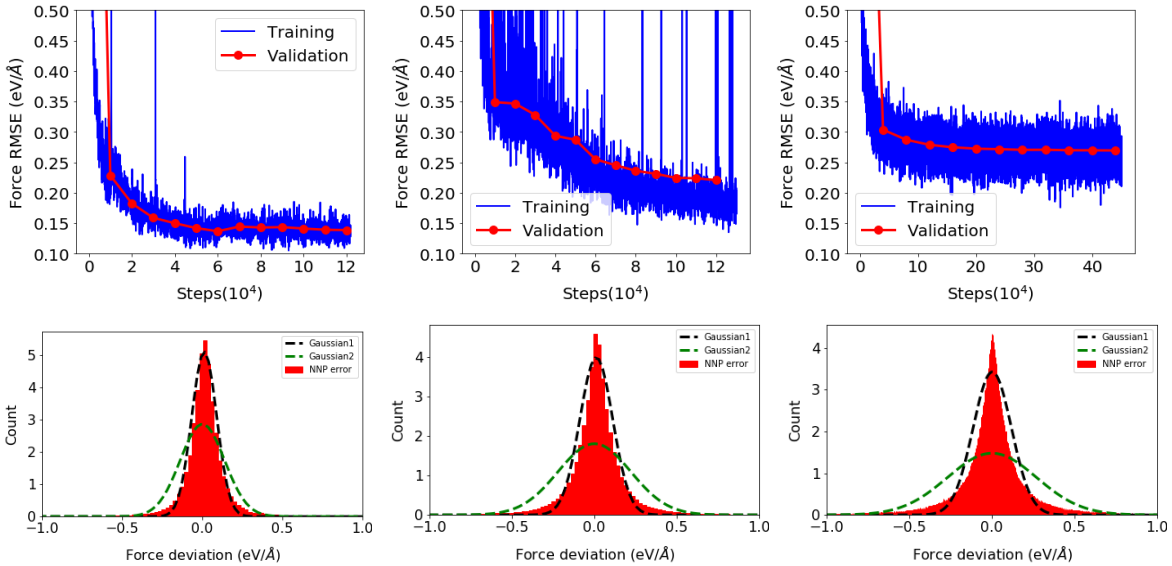

Figure 5.3: Training and validation error in forces. **First row**: The evolution of error in forces on training and validation set for potentials trained at first (*left*), second (*middle*) and third iteration (*right*) of the self-consistent cycle. The blue lines are the RMSE on a given batch of 128 configurations during training. The networks are evaluated as in Fig. 5.2 and on the same number of configurations.(red dots with lines as guide to eye). **Second row**: Error distribution for the validation dataset at first, second and third iteration from left to right for energies (top) and forces (bottom). The histogram is obtained with 3000, 5200 and 12000 configurations respectively. The histogram is constructed such the area is equal to unity. The black dashed line is a normalized Gaussian fit, resulting in a mean $\mu = 0.012$ eV/Å, $\mu = 0.011$ eV/Å, $\mu = 0.002$ eV/Å and standard deviation of $\sigma = 0.08$ eV/Å, $\sigma = 0.10$ eV/Å, $\sigma = 0.12$ eV/Å respectively clearly failing to fit the fat tailed distribution. The green dashed lines are the Gaussian distributions obtained with the zero mean and standard deviation per atom ($\sigma = 0.14$ eV/Å, $\sigma = 0.22$ eV/Å, $\sigma = 0.27$ eV/Å) of the validation dataset, respectively.

5.3.2 Structural and Elastic Properties

In this section, we discuss the performance of the NNP on the structural and elastic properties of selected Carbon polymorphs, namely, diamond, graphite and graphene (See Fig.5.5 and Table 5.2). In Fig.5.5 we show the energy-volume relationship of the three well known polymorphs of carbon. There is excellent agreement of the NNP with the reference DFT results. Next, we report in Table 5.2 the prediction of the elastic properties of these polymorps. The reported equilibrium lattice parameters are obtained by minimizing the total energy until the force components on each atom are lower than 26 meV/Å for both DFT and NNP simulations. We also include results obtained with Tersoff potential, as well as other DFT and machine learning studies in the literature.

In the case of diamond, all machine learning methods agree reasonably well with the DFT results they were trained with, both for the equilibrium volume and elastic constants. The largest deviation is seen in C_{12} prediction in diamond with GAP of Ref. [91] with 24% relative error. For all properties tested, the predictions of NNP of the current study is within a relative error of 5% with respect to DFT. It should be noted that the variation between DFT studies employing different exchange correla-

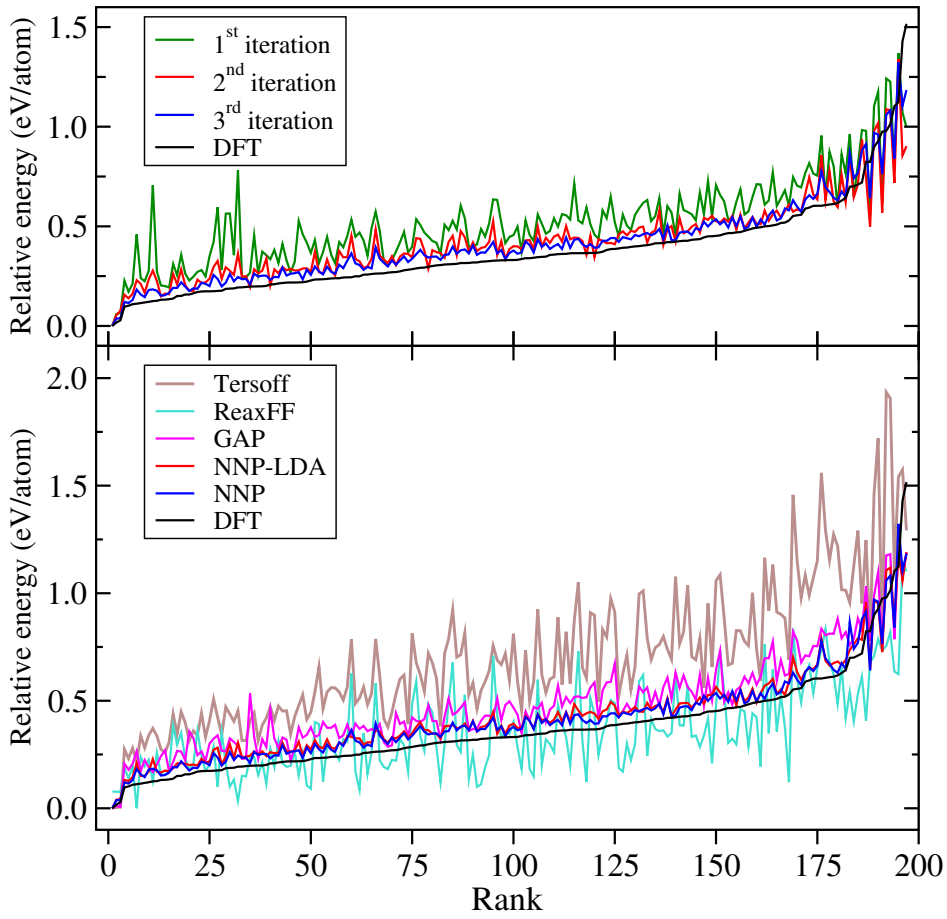

Figure 5.4: Energy ordering of the 197 distinct carbon structure reported in Ref. [130]. **Top:** Performance of NNP at different iterations of the self-consistent cycle. **Bottom:** Performance of GAP [91], reactive force field (ReaxFF) [131] and Tersoff [106] models compared to the final NNP model (*blue line*). For comparison we train a new model with LDA exchange correlation functional, named as NNP-LDA (*red line*).

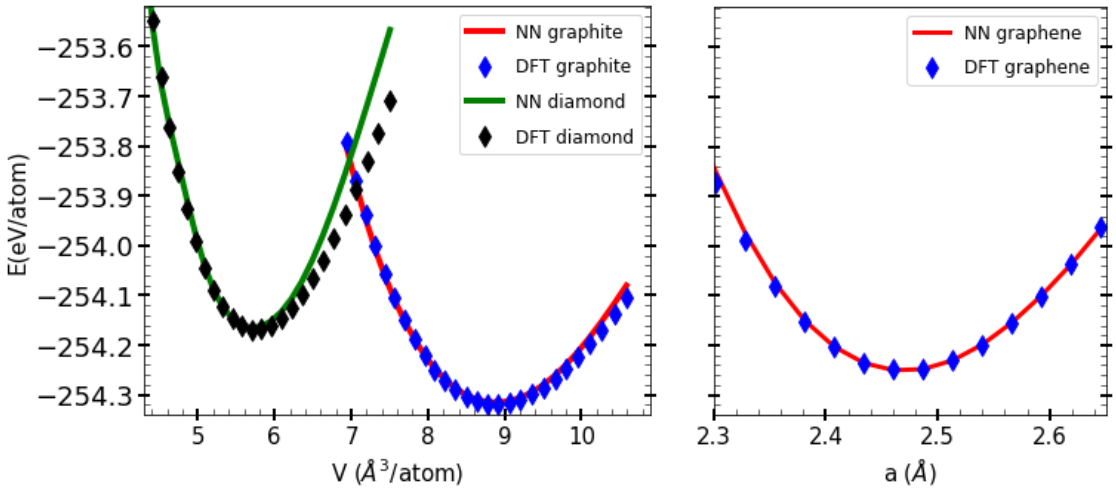

Figure 5.5: Energy as a function of volume with isotropic deformation for diamond, graphite and graphene, computed with the final NNP model and DFT-rVV10.

tion functionals are larger than the difference between machine learnt models and the DFT results they are trained to reproduce. Tersoff potential, although it predicts the

equilibrium volume well, fails to predict the C_{44}.

In the more challenging case of graphite, C_{11} and C_{12} relate to the in-plane elastic properties while C_{33} probes the relationship between strain and stress between the planes which are held together by van der Waals interactions. C_{13} and C_{44} couple the strong in-plane interaction with the weak out-of-plane ones, namely C_{13} can be seen as a measure of interlayer dilation upon layer compression, and C_{44} as a measure of response to shear deformation. The performance of the NNP on prediction of graphite elastic constants are aligned with this overview: For all potentials reported in the middle panel of Table 5.2, in-plane lattice parameter and elastic constants are better predicted than the ones that relate to out of plane interaction, indicating that more data or better training is needed to describe these more delicate properties. Yet it is encouraging that the general purpose NNP of the current work performs at least as well as other NNPs from literature that were developed with a focus on van der Waals systems such as graphite and multilayer graphene. In section 5.3.5 we discuss how focusing on particular system could further improve on these predictions.

5.3.3 Vibrational properties

Phonon dispersion relations give a complete picture of the elastic properties of a material, and reproduction of the dispersion relations obtained via DFT is a tight accuracy criterion on model potentials. Here we examine the performance of NNP through its prediction of phonon dispersion in the case of diamond and graphene, as a function of lattice parameter, up to a 1% deviation from the equilibrium structure. This is a relevant range for thermal expansion of these materials as, for instance, the change in lattice parameter of diamond at temperatures up to 2000 K is found to be below 1 % [141]. Similarly, thermal expansion increases graphene lattice parameter only within 1% at temperatures up to 2500 K [142].

The predictions of NNP for phonon dispersion of diamond and graphene are depicted in Fig. 5.6. There is an overall good agreement between NNP and DFT in the case of diamond. In the case of graphene, there is a slight disagreement for the transverse optical mode around K point. This is the same trend observed in other machine learnt potentials [77, 110] and likely the result of electronic structural properties associated with this special point coupling with the lattice vibration. For both structures, the predicted phonon frequencies reduce when the crystal expands and increase when it is compressed, as expected. An exception to this is the soft flexural mode of graphene close to Γ point. The instability of graphene upon compression can be seen via small imaginary frequency of this mode (shown as negative). This feature is predicted with DFT and is successfully reproduced with NNP, pointing at the capacity of NNP in predicting important structural stability indicators.

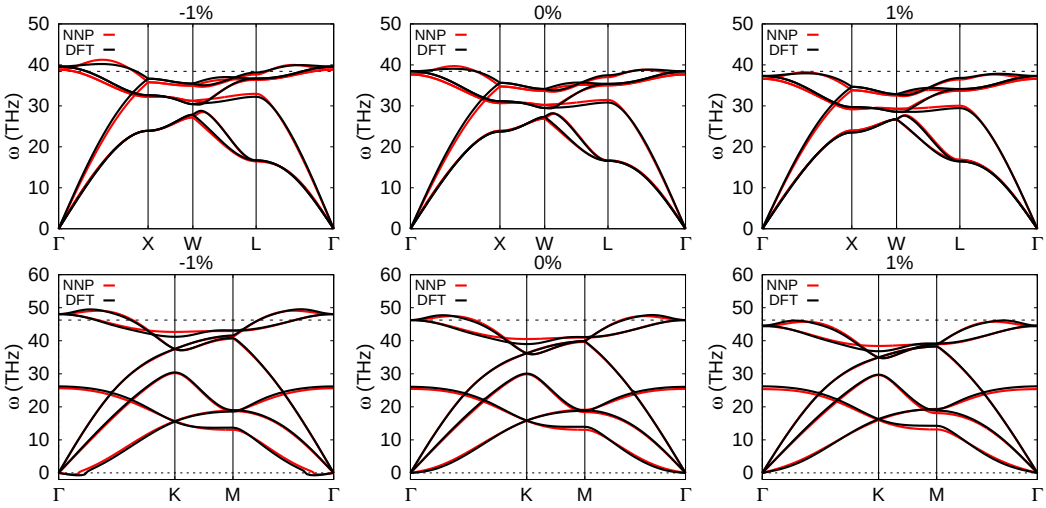

Figure 5.6: Phonon dispersion of diamond (**top**) and graphene (**bottom**) along the high symmetry lines. The value at the top of each graph represent the percentage of expansion (positive) or compression (negative) of the lattice parameter. The black dotted line is the maximum frequency in THz at the Γ point at equilibrium lattice parameter.

Phonon dispersion of graphite, shown in Fig. 5.7 displays negative frequencies for low wave vectors close to Γ, along the perpendicular direction to the graphene plane. These phonon modes are particularly soft and are very sensitive to the level of accuracy of the forces predicted by NNP. We verify this hypothesis with an alternative loss function for NNP training, one that minimizes the *relative* force error rather than the *absolute* one used so far (See Methods section). With a loss function that is based on relative error, configurations with small forces impact the NNP parameter optimization more strongly. We retrain the NNP starting from the previously optimized parameters and report graphite phonon dispersion obtained with the retrained NNP in Fig. 5.7 (middle panel). It is evident that this approach can improve the NNP prediction for structures with small forces, e.g. close to equilibrium conditions. Phonon dispersions for diamond and graphene obtained with this NNP are given in Fig.5.8, and demonstrate that the general quality of the NNP is slightly modified and mostly for the high frequency modes. Further tuning of retraining parameters and loss function can be used as a way to achieve higher accuracy in the desired range of energy and force distributions.

An alternative approach that is commonly used in literature for improving NNP prediction is to bias the training set with the configurations for a certain polymorph. To show the effect of this approach, we train the NNP model from scratch this time using a biased dataset with structures from the close neighborhoods of diamond and/or graphite only. The results reported in Fig. 5.7 (bottom panel) show that this approach indeed allows to reach a better agreement with DFT and there are no imaginary phonon frequencies. However, as it will be further examined later (see Sec. 5.3.5), while this NNP model predicts well properties of configurations around its reference, i.e.

diamond and/or graphite, it is found to be highly non-transferable to other regions of the potential energy surface of carbon.

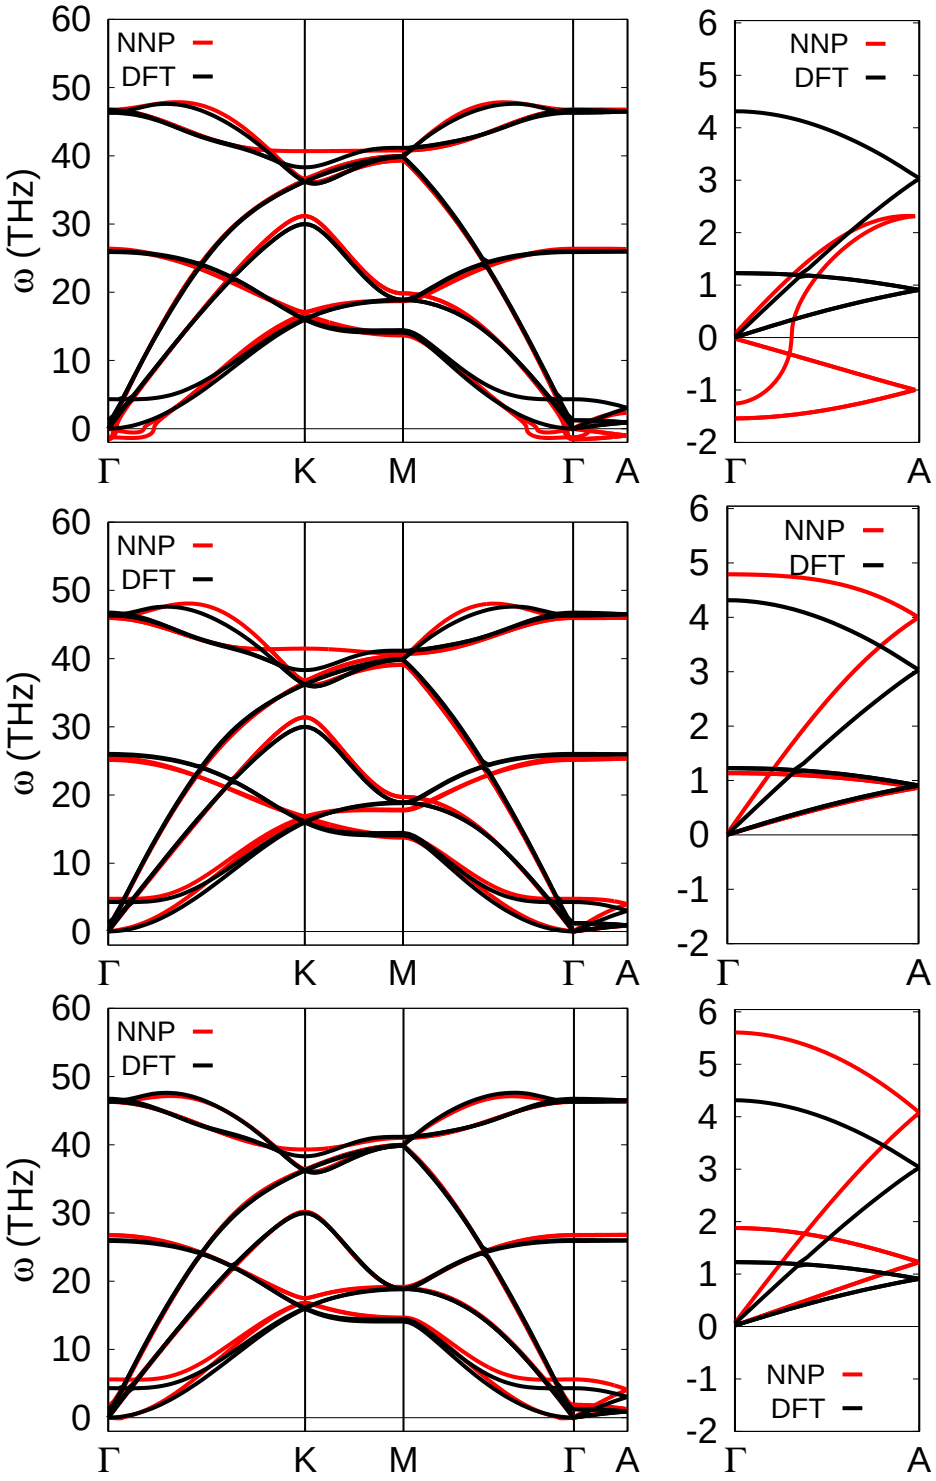

Figure 5.7: Phonon dispersion of graphite along the high symmetry lines for (**top**) an NNP trained with the whole dataset at the last iteration, (**middle**) an NNP retrained with the whole dataset but with the minimization of the relative error on forces and (**bottom**) an NNP trained with all the data within $D = 0.05$ from diamond and graphite (D_{12}, as described in Sec. 5.3.5). The (small) imaginary frequencies are lifted by modifying the NNP training loss function, or by training on data close to graphite in structure.

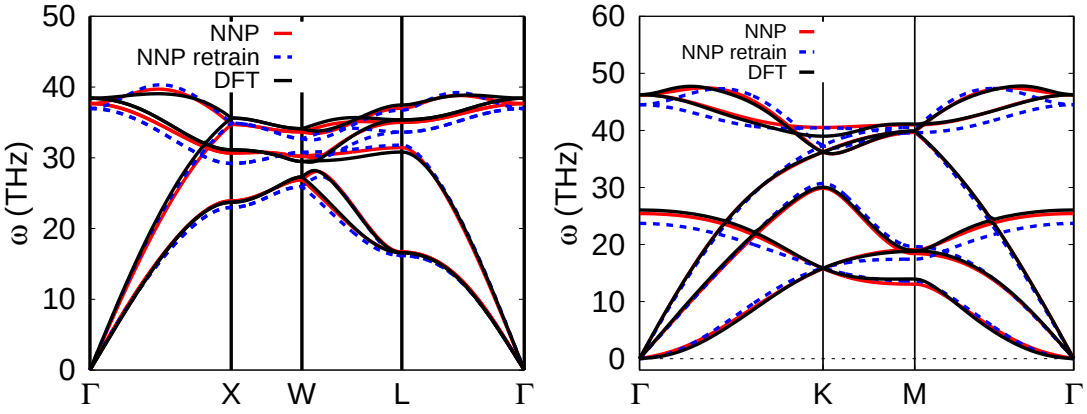

Figure 5.8: Phonon dispersion of diamond (**left**) and graphene (**right**) obtained with an NNP retrained with the updated loss function. The lattice constants upon retrain are unchanged, 3.576 Å for diamond and 2.471 Å for graphene. The phonon dispersion is slightly modified, and mostly for the higher frequency modes as can be expected.

5.3.4 Amorphous carbon structures

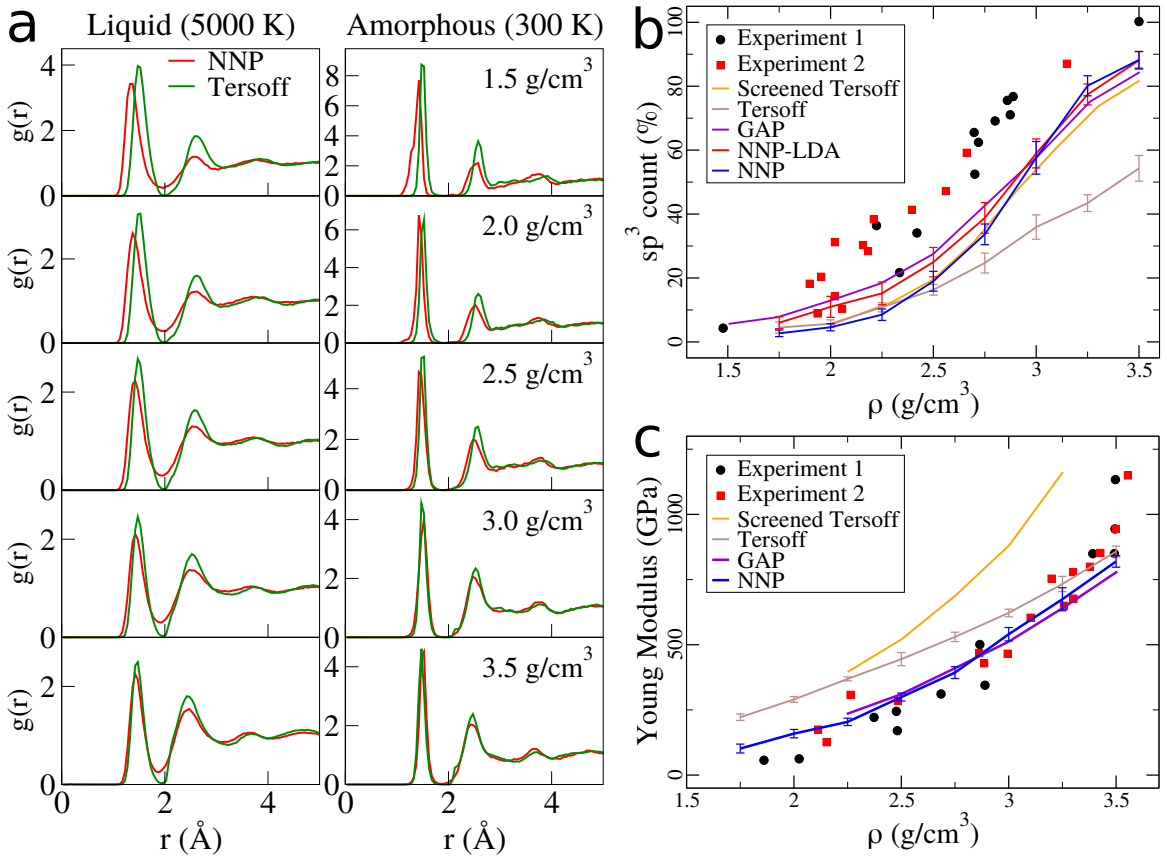

Figure 5.9: The amorphous phase of carbon: (**a**) Radial distribution function for liquid (left) and amorphous (right) carbon, for our NNP and Tersoff potential, at increasing densities (top to bottom). (**b**) Percentage of tetrahedrally coordinated atoms in amorphous carbon structures as a function of density, comparing NNP with rVV10 and LDA-level, and Tersoff potential to results taken from Ref. [91] for GAP and screened Tersoff potentials, as well as experimental results from Refs. [143] and [144]. (**c**) Young Modulus of amorphous carbon as a function of density for NNP at rVV10 level and Tersoff, compared to results taken from Ref. [91] for GAP and screened Tersoff potentials, as well as experimental results from Refs. [145] and [146]

Furthermore, we test the NNP in its ability to construct amorphous carbon structures in a range of densities from 1.5 to 3.5 g/cm^3 generated via the melt and quench method following the steps highlighted in Ref. [91]. We start from a 216 atoms simple-cubic simulation cell and randomized velocities at 9000 K and perform molecular dynamics simulation first at 9000 K with Nose-Hoover thermostat [147] for 4 picoseconds (ps), followed by another at 5000 K for 4 ps, then a fast exponential quench to 300 K at a rate of 10 K/fs (total duration ∼0.5 ps), and finally for 4 ps we let the system evolve with the thermostat fixed at 300 K.

The radial distribution function (RDF) of liquid and amorphous phases are given in Fig. 5.9**a**. The liquid is less ordered than the amorphous configurations at all densities, for all potentials considered. In Ref. [91], it was shown that both DFT and GAP have a non-zero first minimum for the liquid phase at about 1.9 Å which is not properly described by the screened Tersoff potential [148]. Similarly, the NNP of this work captures the non-zero first minimum in the liquid phase while the original Tersoff potential does not. In the case of the amorphous phase, historically one of the first validation cases for the Tersoff potential, the agreement is overall better.

A more detailed comparison of RDF reported in Ref. [91] and experiments is given in Appendix 5.3.4 and shows NNP can successfully reproduce peak position and width across the densities considered.

In order to quantify the short-range order of amorphous structures, we calculate the sp^3 concentration by computing the fraction of carbon atoms with at least four neighbours within a 1.85 Å radius. In Fig. 5.9**b** we show the behavior of this quantity as a function of density, comparing with the results of Ref. [91] and those obtained with regular and screened Tersoff potentials [148]. All methods underestimate the experimental observations yet show a similar general trend with density.

There are quantitative differences among the predictions of theoretical models, in particular, the difference between NNP and GAP predictions are more significant at medium and low densities. This may be attributed to the fact that the DFT dataset used to construct the GAP potential is built with local density approximation (LDA), while in this study the DFT dataset for NNP is built with an accurate exchange-correlation functional that includes van der Waals (vdW) interaction from first principles. In the low density region, vdW interactions allow bonding beyond the typical sp^3 bond length, such that low energy configurations can be constructed with less sp^3 and more sp^2 bonds; while at high densities and at shorter length scales vdW interactions are of lesser significance. This is more evident as we compare the sp^3 count predicted with NNP-LDA as it agrees more closely with the GAP result, revealing the role of the underlying DFT reference in the prediction of the properties of amorphous materials with machine learnt potential models.

The bonding character between atoms strongly affects the elastic properties of ma-

terials. Hence, comparing the elastic properties as observed by experiments with those predicted by theory is another way of assessing the theoretical prediction of sp^3 count in amorphous structures. In order to do that we first find the metastable configurations closest in the phase space to the amorphous structures examined so far, by further quenching the dynamics from 300 K to 0 K, and then performing geometry relaxation until the force components on atoms are below 1 mRy/bohr at fixed volume. Fig. 5.9**c** shows the Young's modulus of these metastable amorphous structures as a function of density. The agreement with the experiment is remarkable, hinting that the discrepancy in theoretical and experimental sp^3 count seen in Fig. 5.9**b** might stem from an inconsistency in definitions between theory and experiment, i.e. the neighbor count within 1.85 Å used in theory underestimates the experimentally measured value that is obtained via comparison of electron energy-loss spectroscopy (EELS) peak area to graphitized carbon [143, 144].

We emphasize that the NNP was not constructed specifically for the description of amorphous C, nor did it include amorphous or melt structures hand-picked to represent these configurations. Despite this, the self-consistent approach yields an NNP which describes these structures well at all volumes considered, validating successful extrapolation of the potential beyond the training set (See Appendix B.4 for energy analysis of liquid and amorphous structures compared to the training set)

5.3.5 Transferability of Neural Network Potentials

The accuracy of a neural network model is often measured by the distribution of the prediction error on a test dataset, in particular via mean and standard deviation of error. But as is the case with training sets, test sets are also not standardized between studies. Therefore the accuracy of potentials tested on different datasets cannot be compared. Here we study the effect of the training and test sets on the apparent accuracy of networks, and measure the impact of these sets on the transferability of neural network potentials.

For every configuration in a dataset we first define its Euclidean distance from a reference atomic environment (e.g. cubic diamond, graphite). The distance between the reference configuration α and a given configuration β is defined as

$$d_{\alpha\beta} = \frac{1}{2}\left(\frac{1}{N_\beta^{\mathrm{at}}}\sum_{i=1}^{N_\beta^{\mathrm{at}}}|\mathbf{g}_\alpha - \mathbf{g}_\beta^i|^2\right)^{1/2} \tag{5.8}$$

where $\mathbf{g} = \frac{\mathbf{G}}{|\mathbf{G}|}$ with $\mathbf{G}$ being a "fingerprint" vector that describes the atomic environment of all atoms in the unit cell for a given configuration, N_β^{at} is the number of atoms in configuration β. In this work, for the definition of atomic environment, we use the

well-established atom-centered symmetry functions of Behler and Parrinello [79], with modifications by Ref. [97] and [75]. This definition is also used to describe the input to the neural network architecture.

Then, we construct a dataset by considering only configurations within a given cutoff distance D from this reference. Following this strategy we build four datasets, three of which are referenced from cubic diamond with D values of 0.05, 0.10 and 0.15; the fourth one is referenced from either cubic diamond or graphite with D=0.05 (denoted by D_{12}). For each D, 20 % of the dataset is set aside for validation and the remaining 80 % is used for training. We train four different NNPs on these four sets from scratch, and test each on the respective validation datasets.

In the top panel of Fig. 5.10, we report the training and validation RMSE in energy prediction as the cutoff distance D from the reference structure increases. We show that an RMSE as low as 2.4 (2.5) meV/atom for training (validation) can be obtained when training and validation configurations are very similar, i.e. within a distance of 0.05 from the diamond reference. However, the prediction error of this NNP dramatically increases as it gets tested on structures farther in the input space, to as high as an RMSE of 473 meV/atom. This is a confirmation of the common observation that the prediction error of a neural network is strongly dependent on the similarity of training and test environments [149]. On the other hand, when the model is trained and tested using the complete set, a prediction RMSE of 22.1 meV/atom is obtained for energy, while, for the configurations within $D = 0.05$ from diamond, the prediction RMSE is still considerably small, 7.7 meV/atom. The analysis for forces follows the same trend as energies. The RMSE values for energies and forces are given in Appendix B.5.

Hence, it can be deduced that, for a fixed network architecture, a trade-off must be struck between having small error on configurations similar to a reference structure, and obtaining reliable predictions for general configurations from the full potential energy surface. The other entries in these tables confirm this analysis: the more diverse the training set is, the more robust is the resulting potential outside its training basin. Therefore, for a reliable NNP for multiple C polymorphs, as the one targeted here, a diverse training set from a wide region of the potential energy surface is necessary.

5.3.6 Conclusion

In summary, we have presented a self-consistent technique for generating an accurate and transferable neural network potential. Since neural networks encode the physics of a system into their parametrization through data, the dataset plays a crucial role in the resulting NNP performance. The method described in this work achieves a comprehensive dataset via balanced integration of evolutionary algorithm, unsupervised machine learning in the form of clustering, and molecular dynamics. As the training

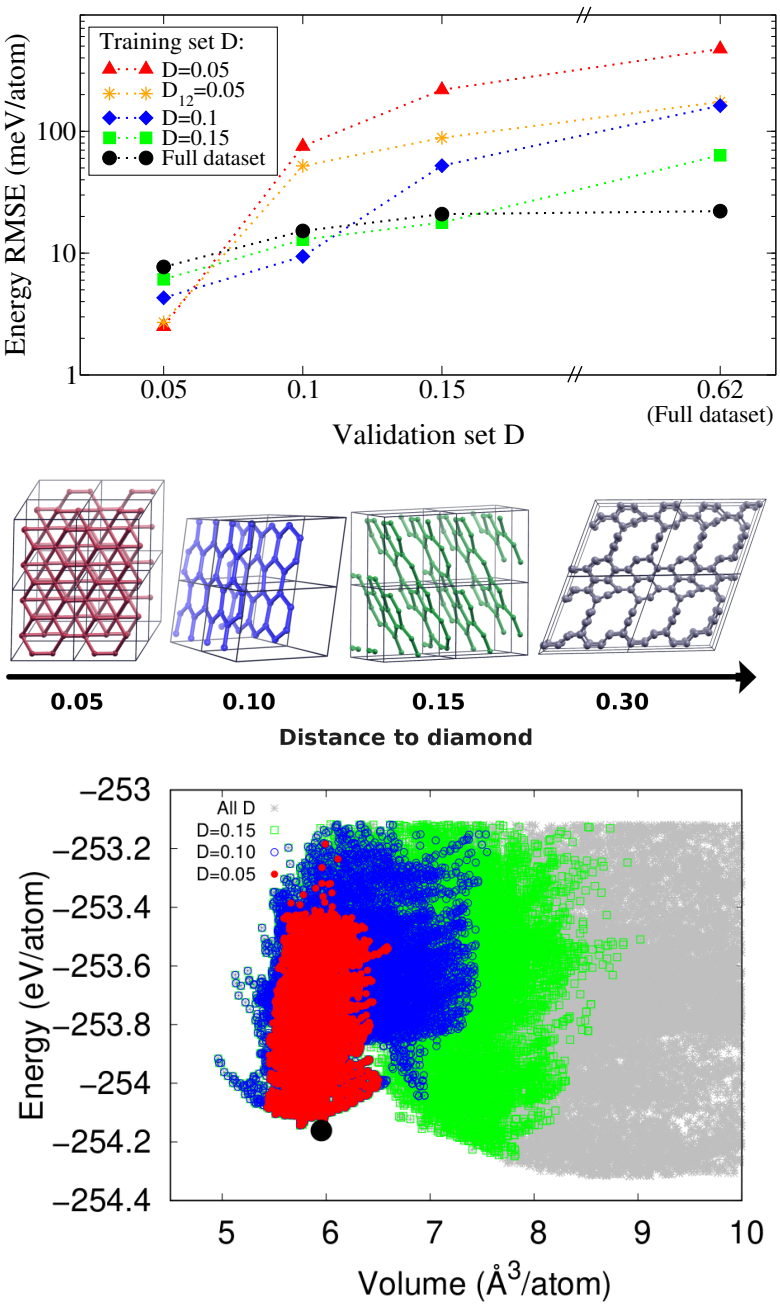

Figure 5.10: **Top:** Validation error of networks trained on different datasets as a function of the distance of the validation set from diamond. Numerical values are given in Appendix B.5 for energies and forces (not shown here). **Middle:** Representative structures at given distances from diamond, the reference structure. The structures at 0.05 or lower are recognizably related to the reference, while at 0.10 and 0.15 compressed and/or defected layered structures are visible. At 0.30 and above, configurations with several double bonds and Carbon chains appear. **Bottom:** Energy per atom as a function of volume for structures in the dataset, colored according to their distance cutoff D from diamond. The black dot corresponds to the reference diamond structure. The complete dataset includes structures with larger volume that are omitted here for clarity. The complete volume range is given in Fig. B.4 in Appendix B.5

dataset is central to all machine learning models, we believe this generation method may be adopted by and would be beneficial to other ML approaches as well.

The distance-based analysis also gives an *a posteriori* measure of the profound

diversity of the final dataset achieved via the self-consistent method. MD together with EA and clustering successfully explores a wide range of configurations on equal footing so that the dataset shown in Fig. 5.10 covers energy and volume landscape rather homogeneously. This is in line with the observation that at each iteration dataset diversity increases and validation RMSE may also increase since the network is tasked with a more complex functional approximation problem.

The presented workflow requires minimum human intervention. As the potential is iteratively improved, even rough starting models could be utilized for the very first step, and we have shown that the converged potential does not carry the limitations of the initial model. Therefore, not only this workflow is ready for high-throughput automation schemes as envisioned in future of experimentation but it is also robust with respect to lack of previous information about a system, as is often the case with novel materials.

Many new materials with practical applications can be expected to be multicomponent systems. As the phase space of possible compounds grows larger and wildly unexplored, truly automated and unbiased approaches for an efficient exploration will become essential. We believe that our dataset generation approach (which can be coupled to any other ML approximator with multicomponent capability, e.g. Ref. [150]) would be particularly suited to such systems. The workflow and the underlying neural network [98] and electronic structure codes are publicly available and are open-source.

The self-consistent NNP generation procedure is entirely system independent and we demonstrated its successful application to the challenging case of Carbon for which classical and machine-learnt potentials are abundant in literature. We show that for diamond, graphite and graphene phases, NNP reported in this work performs considerably better than Tersoff, a classical potential, and overall better than the existing machine learnt potentials for structural and elastic properties. Recently, a new GAP model trained on a large dataset with wide range of polymorphs was published [151].

When predicting graphite phonon dispersion, NNP resulted in very good agreement for the majority of the modes, yet predicted instability for the very soft modes that relate to interlayer interaction. We have traced this behaviour to the accuracy requirement in predicting such small forces. To increase accuracy using a fixed neural network architecture, we built the training set only with structures that are in the vicinity of graphite according to a fingerprint based distance measure. The resulting potential provided accurate phonon frequencies but it showed poor generalization to a wider range of structures, compared to a more comprehensive potential trained on the entire dataset. This example highlights the need for a procedure to standardize the accuracy measure of NNPs and a more pressing need to build error estimate measures into the process of generating NNPs.

Chapter 6

Extension of PANNA to include Long Range Electrostatic Interactions: the case of Li-C systems

6.1 Introduction

The limitations of the cluster expansion methods highlighted in chapter 4 prevent the construction of a site-based interatomic potentials for Li interaction with a generic carbon host. The experimentally observed high-lithium capacity carbon-based anode materials are three-dimensional (3D), highly-porous carbon framework composed of two-dimensional (2D) graphene nanoflakes with a high concentration of defects, vacancies and edges [5, 152]. These materials are most often modelled in theoretical studies with pure and defective single-layer graphene [8–16, 53] as well as pure few-layer graphene [8] because of the high computational cost of performing both zero and finite temperatures *ab initio* calculations. In chapter 3 we presented a simplified model that enables a detailed study, for the first time, of the finite temperature effects of Li interactions with graphene. In this study, Li-cluster was found to be stable on graphene substrate at high concentration, an undesirable phase of Li in aqueous environment that may lead to internal short-circuit of the battery when the graphene technology is eventually used as the anode of the lithium ion batteries (LIBs). However, considering the production mechanisms of these high Li density anode carbon-based materials (See Ref. [5] for example), a combination of various effects such as bending, defects, edges, vacancies etc., are expected to be present and may enhance the stability of disperse arrangements of Li on these materials even at room temperature. In order to study these effects at experimental conditions via atomistic simulations, a general Li-C interatomic potential with *ab initio* accuracy is necessary.

We are aware of two Li-C interatomic potentials available in literature. The first one

is the Li-C interatomic potential constructed within the reactive force-fields (ReaxFF) methods [131] whose parameters were optimized to reproduce *ab initio* elastic constants and cohesive energies of Li adatom and small Li-clusters on pure graphene and graphite configurations. This potential, although constructed to reproduce *ab initio* data, is less general and hence is not suitable for the current study. In fact, we have demonstrated in chapter 5, that this ReaxFF potential perform poorly in ranking pure C configurations (see Fig. 5.4).

The second potential is the recently developed Gaussian approximation potential (GAP) [153] for Li-C systems fitted to reproduce a wide range of Li-C configurations. The training dataset was generated by randomly inserting Li on randomly distorted 24-atom cells of graphene and graphite as well as 64-atom cell of amorphous carbon obtained from melt and quench procedure. An effective two-body Li-Li potential was added to the standard GAP model to account for long-range interactions. Although the GAP potential has been demonstrated to be very accurate [96] and in principle could describe Li interaction with generic C environments as we demonstrated in chapter 5, machine learning model fitted on limited dataset suffer from transferability problems.

In this chapter, we present our preliminary results on the construction of a general Li-C interaction potential within artificial neural network (ANN) machine learning formalism following the self-consistent approach proposed in chapter 5. Proper treatments of electrostatic interaction is incorporated in the existing short-range ANN potentials implemented in PANNA package via charge equilibration techniques as will be discussed in Sec. 6.2.1. This techniques of combining standard ANN potential with long-range electrostatic interaction via charge equilibration is new. A similar approach had just appeared on arXiv [154] at the time of the writing of this thesis. The similarity and differences between our current implementation and the one in Ref. [154] will be highlighted in Sec. 6.2.1.

6.2 Methods

6.2.1 Artificial Neural Network with Long-Range: Theoretical consideration

The artificial neural network potential (ANNP) based on atom-centered descriptors [79] neglects long-range interactions beyond a spherical cutoff, R_c. The interaction in systems mediated by charge transfer (for instance, ionic systems such as sodium-chloride or Li adsorption in C networks) that decays as $1/r$ are neglected. Likewise, systems with significant van der Waals interactions due to non-local fluctuation in density (for instance, molecular crystals or graphite) are also not properly described in the cutoff atom-centered scheme. The van der Waals interactions, although very

important for the stability of some systems, are neglected in the current study. Instead, the long-range electrostatics interactions are the main contributions that stabilize Li-C systems as a result of charge transfer from Li to the carbon host.

The absence of long-range electrostatics interactions in ANNP was identified few years after Behler and Parrinello-proposed atom decomposition, that permits the application of ML-based potentials to extended systems, and efforts have been made to address the problem. One of such efforts dated back to 2012 when Morawietz et al. [155] constructed interatomic potential for water dimer in which, in addition to the short-range atomic energies, a long-range electrostatics interactions was incorporated via environment dependent atomic charges. Separate neural network models were constructed for the atomic charges and the parameters of the model were optimized to reproduce *ab initio* charges obtained from Hirshfeld charge decomposition scheme [35]. In this scheme, the total energy was trained in two steps: (i) the electrostatic Coulumb energy was computed from the environment dependent charges trained on Hirshfeld charges and removed from the total energy, and (ii) the remaining energy assumed to be short-ranged and trained using short-ranged atomic energies.

Another approach to include long-range electrostatic interactions in artificial neural network interatomic potentials is through the charge equilibration scheme proposed by Ghasemi et al. [156]. Although this method was developed to construct a general ANNP including short-range contributions, it has mostly been successfully applied to construct ANNP for ionic systems [156–158]. Its main idea is based on the fact that the environment dependent atomic charges can be obtained from short-ranged environment dependent atomic electronegativities through a charge equilibration scheme [108]. By construction, it allows non-local charge redistribution in contrast to the scheme proposed in Ref. [155] in which charges are a function of the local environment only. The total energy expression is similar to the standard total energy in charge equilibration scheme (see Ref. [108]) except that the point charge Coulumb energy is replaced with the Hartree energy of a Gaussian charge density distribution. The procedures to predict the total energy are as follows: (i) construct a neural network for the environment dependent electronegativity, χ, (ii) solve a linear systems to obtain the atomic charges, q. The linear system is a result of the minimization of the total energy subject to charge neutrality conditions, and (iii) compute the total energy from the χ and q.

In this work, we adopt the charge equilibration scheme for the description of the long-range electrostatic interaction because of the following advantages :

1 it does not require the decomposition of the *ab initio* charge density into atomic components, although the charges can be constrained during the training.

2 charge neutrality is satisfied by construction.

3 long range charge transfer is permitted.

The total energy of a given configuration is decomposed into two parts: the short ranged part that is obtained from the standard atomic energy contributions and the long-range part obtained from the charge equilibration scheme explained above. Thus, the total energy then reads as:

$$E_{tot} = \eta_{SR} E_{SR} + \eta_{LR} E_{LR} \tag{6.1}$$

where E_{SR} is the original total energy which is the sum of environment dependent atomic energies, $E_i(G(\{\mathbf{r_i}\}))$, η_{SR} and η_{LR} are the fraction of the short range and long range energy contribution per cell, E_{LR}, defined as:

$$\begin{aligned} E_{LR}[\rho] = \sum_i \left[E_i^0 + \chi_i q_i + \frac{1}{2} J_i q_i^2 \right] \\ + \frac{1}{2N_{cell}} \sum_{\mathbf{R}} \int \frac{\rho(\mathbf{r})\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}' + \mathbf{R}|} d^3r d^3r'. \end{aligned} \tag{6.2}$$

The energy term, E_{LR}, comprises of atom specific terms and a Hartree term (written for periodic systems) to describe an effective interaction between charge density at $\mathbf{r}$ and $\mathbf{r}'$. χ_i is an environment dependent atomic electronegativity and J_i is the species dependent atomic hardness, which for simplicity, is assumed to be independent of atomic local environment. The sum over R runs through all unit cells in the periodic solid with N_{cell} unit cells.

The total charge density is obtained from the superposition of spherical Gaussian functions centered around each atom i at position $\mathbf{r}_i$ normalized such that the charge density integrates to the total number of charge in the system. The atom-centered charge density is thus defined as:

$$\rho_i(\mathbf{r}) = \frac{q_i}{\alpha_i^3 \pi^{\frac{3}{2}}} \exp(-\frac{|\mathbf{r} - \mathbf{r_i}|^2}{\alpha_i^2}), \tag{6.3}$$

and the total density at position r is then given by:

$$\rho(\mathbf{r}) = \sum_{i=1}^{N} \rho_i(\mathbf{r}) \tag{6.4}$$

where q_i is the charge on atom i and α_i is the Gaussian width of atom i that defines the spatial extent of the charge density around atom i.

In aperiodic systems, the Hartree energy can be evaluated analytically in real space.

Performing the integration gives rise to

$$E_{LR} = \sum_i \left[\chi_i q_i + \frac{1}{2}(J_i + K\frac{2\gamma_{ii}}{\sqrt{\pi}})q_i^2\right] + \frac{K}{2}\sum_{i,j,i\neq j} q_i q_j \frac{\mathrm{erf}(\gamma_{ij}|\mathbf{r}_i - \mathbf{r}_j|)}{|\mathbf{r}_i - \mathbf{r}_j|} \tag{6.5}$$

where $K = \frac{1}{4\pi\epsilon_0} = 14.39964$ eVÅ and $\gamma_{ij} = \frac{1}{\sqrt{\alpha_i^2+\alpha_j^2}}$

For periodic systems instead, the Hartree energy term in this work is evaluated in reciprocal space. The long range term then reads as:

$$E_{LR} = \sum_i \left[\chi_i q_i + \frac{1}{2}J_i q_i^2\right] + \frac{K}{2}\sum_{i,j} q_i q_j \frac{4\pi}{\Omega}\sum_{G\neq 0}\frac{\exp(i\mathbf{G}.\mathbf{r}_{ij})}{G^2}\exp(-\frac{G^2}{4\gamma_{ij}^2}) \tag{6.6}$$

Writing the E_{LR} in compact form we have

$$E_{LR} = \sum_i \left[\chi_i q_i + \frac{1}{2}J_i q_i^2\right] + \frac{K}{2}\sum_{i,j} q_i q_j V_{ij} \tag{6.7}$$

where

$$V_{ij} = \begin{cases} \frac{\mathrm{erf}(\gamma_{ij}|\mathbf{r_i}-\mathbf{r_j}|)}{|\mathbf{r_i}-\mathbf{r_j}|}(1-\delta_{ij}) + \frac{2\gamma_{ij}}{\sqrt{\pi}}\delta_{ij} & \text{No PBC,} \\ \frac{4\pi}{\Omega}\sum_{G\neq 0}\frac{\exp(i\mathbf{G}.\mathbf{r}_{ij})}{G^2}\exp(-\frac{G^2}{4\gamma_{ij}^2}) & \text{PBC.} \end{cases} \tag{6.8}$$

In order to obtain the atomic charges, we minimize the total energy with respect to q_i, subject to the charge conservation rule $\sum_i q_i = q_{tol}$ through a Lagrange multiplier, that is;

$$\frac{\partial}{\partial q_i}\left[E_{tot} - \lambda(q_{tot} - \sum_i q_i)\right] = 0. \tag{6.9}$$

This gives rise to

$$-\chi_i = \sum_j q_j A_{ij} + \lambda, \tag{6.10}$$

where

$$A_{ij} = \begin{cases} V_{ij} & \text{if } i \neq j \\ J_i + V_{ii} & \text{if } i = j. \end{cases} \tag{6.11}$$

Given the atomic electronegativity, χ_i, equation 6.10 combined with the charge neutrality condition, $\sum_i q_i = q_{tol}$, can be solved to compute the atomic charges for a certain choice of J_i and the atomic Gaussian width, α_i.

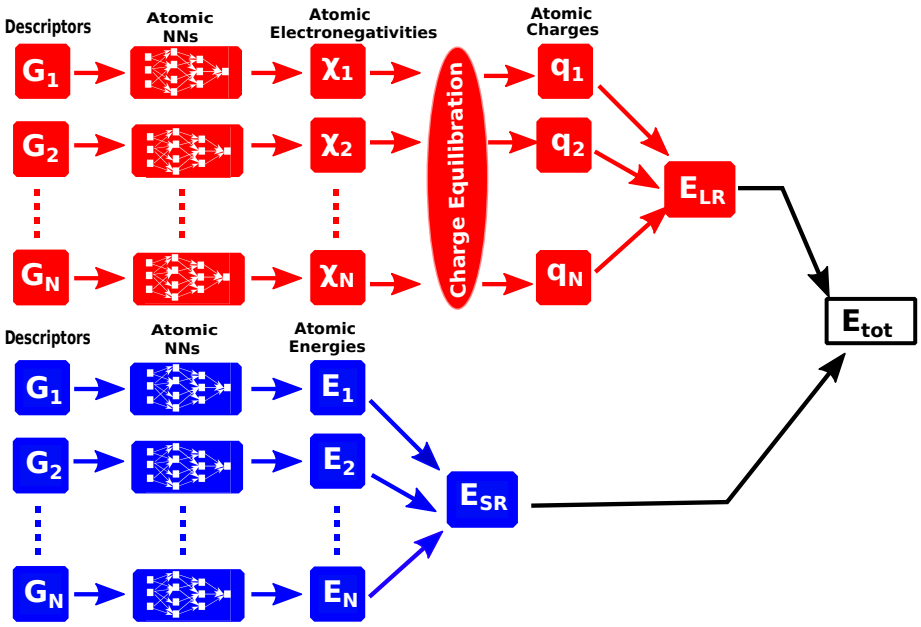

Figure 6.1: PANNA long-range workflow: The red block represent the workflow for the prediction of atomic electronegativities, computation of charges via charge equilibration and the computation of the long-range contribution to the total energy. The blue block represent the prediction of atomic energies and computetaion of the short range contribution to the total energy. The final output is the sum of the two contributions. The input layer represented by G_i $(i = 1, .., N)$ are list of numbers denoting the descriptor of the chemical environment of atom i. N is the number of atoms of the configuration under consideration.

In our implementation, the total energy is computed as follows (see Fig. 6.1 for illustration):

1 Predict the atomic electronegativity from neural network for each atomic environment descriptor

2 Compute the atomic charges by solving the linear system in equation 6.10.

3 Compute E_{LR} by substituting the predicted electronegativity and the computed charges into equation 6.7.

4 Predict the environment dependent atomic energy, $E_i(G(\{\mathbf{r}_i\}))$.

5 Finally, the predicted total energy E_{tot} is then obtained via equation 6.1.

The total force on atom k in the direction β is given by

$$f_{k\beta} = -\frac{\partial E_{tot}}{\partial r_{k\beta}} = -\frac{\partial E_{SR}}{\partial r_{k\beta}} - \frac{\partial E_{LR}}{\partial r_{k\beta}}. \tag{6.12}$$

The contribution of the short range energy term to the total force is computed via chain rule as discussed in chapter 4. The long range contribution to total atomic forces

contains two terms as follows

$$f_{k\beta}^{LR} = -\sum_i q_i \frac{\partial \chi_i}{\partial r_{k\beta}} + \sum_i q_i q_k F_{ik\beta}, \tag{6.13}$$

$$F_{ik\beta} = \begin{cases} \left[\frac{\mathrm{erf}(\gamma_{ik} r_{ik})}{r_{ik}^3} - \frac{2\gamma_{ik}}{\sqrt{2} r_{ik}^2} \exp(-\gamma_{ik}^2 r_{ik}^2)\right] r_{ik\beta} & \text{No PBC,} \\ \frac{4\pi}{\Omega} \sum_{G \neq 0} \frac{\sin(\mathbf{G}.\mathbf{r}_{ik})}{G^2} \exp(-\frac{G^2}{4\gamma_{ik}^2}) G_{k\beta} & \text{PBC.} \end{cases} \tag{6.14}$$

The derivative of χ with respect to r in equation 6.12 is computed via chain rule similar to the short-range part while the part of the last term of the equation that depends on atoms coordinates is precomputed and the sum performed once the charges are predicted.

In the current implementation, the short range energy contribution is assumed to be independent of atomic charges and the dependence of forces on local redistribution of charges is obtained via the change in the atomic electronegativity with respect to atomic coordinates. Instead, in Ref. [154], the short range contribution depends on the atomic charges and therefore, the atomic forces depend on the derivatives of atomic charges with respect to atomic coordinates. This dependence introduced $3N$ linear systems of size $N+1$ for N atom systems similar to the one in equation 6.10. However, by defining an auxiliary function that introduced atom-dependent Lagrange multiplier that constraint the solution of equation 6.10, only one linear system is solved, thereby reducing the computational cost of solving $3N$ linear systems at each training steps.

As depicted in Fig. 6.1, the parameters of atomic networks for energies and electronegativities are currently trained concurrently. Only the total energy and/or forces are required to train the models. As in the case of training of the short range term in which the model is given the freedom to choose a partition of atomic energies so that the total reference energy is correctly predicted, we also do not imposed any constraint on the partitioning of the energy into long range and short range contributions. This is at variance with Ref. [154] in which the Hirshfeld partitioning of total charge density into atom centered densities is used to impose a constrain on the electrostatic contribution to the total energy.

6.3 Implementation

The workflow shown in Fig. 6.1 is an extension of PANNA to include long-range electrostatics through charge equilibration. The philosophy is the following: two networks are created per species, one for the atomic energy and the other one for atomic electronegativity. The two networks are trained concurrently.

The matrix elements V_{ij} and F_{ij} in equations 6.8 and 6.14 are precomputed and saved into a file alongside the atomic environment descriptors as discussed in chapter

4.

These new features in PANNA do not introduce additional scripts to the ones presented in Sec. 4.3.6 of chapter 4 but include new flags (variables) required to specify the calculation of the long-range terms while building descriptors or training of electronegativity during parameter optimization. In the following, we highlight these new flags. Keywords enclosed in square brackets indicate sections in the "python config" file format

Input Binary

```
[LongRange]
  species_gaussian_width:
      comma separated values of species Gaussian
      width in angstrom. There is no default.
  long_range:
      boolean to decide whether kernels are saved
      Default: False
  acc_factor:
      float to determine the truncation of the sum in
      reciprocal space.
      Default:1e-6
```

We truncate the sum when the next term contributes less than acc_factor to the energy in absolute value. The default is 10^{-6}

Training

By default, only the electronegativity, χ is trained, the charges computed with equation 6.10 and the energy computed using equation 6.7. In order to train both electronegativity and the atomic energies, the factor η_{SR} (named SR_factor in the code) must be greater than zero. The total short range contribution to the total energy is this factor multiplied by the network prediction. For training, no new section is introduced a few flags are however added. The new flags/variables are described below.

```
[DATA_INFORMATION]
    atomic_hardness:
        comma separated values of species hardness in eV.
        Default:None
    scaffold_type=LongRange
    chi_offset:
        comma separated values of initial species
```

```
        electronegativities in eV.
        Default:0 per species

[TRAINING_PARAMETERS]
    SR_factor: float.
             Default: 0 implying that there is no
             short range prediction.
    LR_factor: float.
               Default: 1.0 implying that the
                 electronegativity is predicted.
```

Testing

Once a network is properly trained, the evaluation is done with "numpy" with the same script as the one used for the short-range contribution. However, an addition section has been introduced to instruct the code to evaluate the electronegativity networks. The section is highlighted below.

```
[LongRange]
    atomic_hardness:
       comma separated values of species hardness in eV.
       There is no default.
    long_range:
        True if electronegativity is trained.
    SR_factor: float. The default is 0 implying that there is no
               no short range prediction.
    LR_factor: float. The default is 1.0 implying that the
               electronegativity is predicted.
```

In the current implementation, the linear system in equation 6.10 is solved by inverting the matrix A in both training as well as during evaluation.

LAMMPS Plugin

We have also implemented a plugin to use the optimized model in LAMMPS. This is an extension of the LAMMPS plugin implemented for the short range potential. Here, the energy and forces are evaluated in two parts: in real space and in reciprocal space. The matrix A and the Hartree energy are computed in reciprocal space adapting the Ewald summation scheme in LAMMPS. The linear equation is first decomposed into two linear equations by expanding the charge in powers of the Lagrange multiplier in

equation 6.10, λ, up to the first order as described in Ref. [159] and summarize as follows

$$q_i = s_i - \lambda t_i. \tag{6.15}$$

Substituting equation 6.15 into equation 6.10 and imposing the charge neutrality we obtain the following;

$$\begin{aligned} -\chi_i &= \sum_j A_{ij} s_j, \\ 1 &= \sum_j A_{ij} t_j, \end{aligned} \tag{6.16}$$

and

$$\lambda = \frac{\sum_i s_i - q_{tot}}{\sum_i t_i}$$

.

The linear equation is solved iteratively to compute charges using the preconditioned conjugate gradient optimization techniques [159]. The system is initialized by setting s_i to the atomic charge of atom i in the previous MD/relaxation step. In this way, only a few steps are needed to reach convergence during successive steps of molecular dynamics.

6.4 Training Dataset

As demonstrated in chapter 5, a flexible and transferable ANN potential can be obtained via a self-consistent approach that combines evolutionary algorithm, unsupervised clustering algorithms with training of ANN potential. In the case of adsorption problem, a careful selection of model concentrations of foreign atoms (lithium) in the host carbon system network is desirable. The high concentration part of the phase diagram is of particular interest in this work in order to construct an interatomic potential suitable to simulate the anode materials in experimental conditions. Therefore, we devote special attention to the high Li density region of the phase diagram. In Table 6.1, we show the models we based our data generation on. This wide range of concentrations is believed to span all the important range. The scope of the work is currently restricted to 3D and 2D during the evolutionary algorithm search. The work is not expected to reproduce Li adsorption on one dimensional (1D) (e.g carbon nanotube). We however note that the model can be extended to nanotubes whenever the need arises.

Table 6.1: Model inputs for the evolutionary algorithms (EA): The mark, $\surd$, indicates that the corresponding number of Li, N_{Li}, is coupled with the indicated number of carbon, N_C. For instance, we simulate separately cells of 60 C atoms with 1 Li atom and 60 C atoms with 2 Li atoms during the EA search

N_C	N_{Li}			
3D search				
	1	2	3	4
4	$\surd$			
6	$\surd$	$\surd$		
8	$\surd$	$\surd$		
12	$\surd$	$\surd$	$\surd$	$\surd$
16	$\surd$	$\surd$		
18	$\surd$	$\surd$		
20	$\surd$	$\surd$		
24	$\surd$	$\surd$	$\surd$	
60	$\surd$	$\surd$		
2D search				
8	$\surd$			
12	$\surd$	$\surd$	$\surd$	
24	$\surd$	$\surd$		

6.4.1 Evolutionary Algorithm for Configuration Space Search

The crystal structure search was carried out with the evolutionary algorithm (EA) as implemented in USPEX [119]. As mentioned in chapter 5, the number of iteration steps required to reach self-consistency in our proposed approach depends on the starting interatomic potential. The final interatomic potential however, carries no signature of the initial potential. Here, we start with the general potential developed in chapter 5 for C-C interactions coupled with the ReaxFF used as the starting potential in that chapter to describe Li-Li and Li-C interactions. Since ReaxFF is a many-body interaction potential, we carefully set all the coefficients of C-C interactions of the potential required to set the interaction to zero.

At each stoichiometry tabulated in Table 6.1 (also pure Li-atom cells with 4 and 8 atoms, not shown in the table), we start EA with random structures with population size of $2 \times N_{at}$ for N_{at} atom unit cells. The population size is set to 50 whenever $2 \times N_{at}$ is greater than 50. The EA is evolved with the following evolutionary operations: 50 % of structures are generated by heredity, 20 % by mutation, 10 % by permutation and 20 % are new random configurations.

At each generation, structures are optimized in five successive steps: (a) constant pressure and temperature molecular dynamics at 0.1 GPa and 50 K respectively for 0.3 ps with time step of 0.1 fs, (b) relaxation of cell parameters and internal coordinates until force components are less than 0.26 eV/Å, (c) constant pressure and temperature molecular dynamics at 0.1 GPa and 50 K respectively for 0.3 ps with time step of 0.1 fs,

(d) relaxation of cell parameters and internal coordinates until force components are less than 0.026 eV/Å, and (e) a final relaxation of cell parameters and internal coordinates until force components are less than 0.0026 eV/Å.

The EA algorithm was evolved up to a maximum of 20 generations. This procedure gives rise to a diverse variety of polymorphs.

6.4.2 Clustering

In order to sample around the minimum of each polymorphs within a given stoichiometry (see Table 6.1), we perform , an unsupervised, bottom-up, hierarchical clustering based on distance measure approach. Instead of the single linkage used in chapter 5 in which two distinct clusters with closest pair of elements are agglomerated, here we used the complete linkage or farthest neighbour clustering that form a bigger cluster out of two clusters using the farthest pair as a measure. This latter approach is found to perform better in grouping configurations with distinct environments. For each stoichiometry, we extract configurations in a stable clusters at a cutoff distance of at least 0.1. The distance measure is based on the cosine distance defined in equation 5.2 of chapter 5.

A single configuration was selected from each cluster center at random and classical molecular dynamics was performed for a total time of 50 ps with the Nose-Hoover thermostat and barostat. In these simulations the external conditions of pressure and temperature are ramped up from -10GPa at 100 K, to 10GPa at 1000K in the course of 50 ps. A total of 50 configurations were selected from the 50 ps simulation at an interval of 1 ps.

In order to avoid to construct an unphysical potential for Li-C interaction in graphene as Li is moved away from the surface of graphene, we randomly selected 50 configurations from Li:24C 2D configurations obtained from the evolutionary search and randomly modified the Li position on the 50 configurations to sample long Li-C distances as follows: given the lattice constant of graphene,$a = 2.47$ Å, for each configuration, choose two random numbers in the interval $(0,a/2)$ and add them to the current x and y coordinates of the Li atom and sample the z component in a deterministic manner by adding distances $-0.8, 0.5, 1.0, 2.0, 2.5$ Å to the z component of the Li position. This procedure became necessary since these configurations are high in energy and cannot be reached via classical molecular dynamics.

As we aspire to construct an interatomic potential for Li-C applicable to study also pure Li and pure C configurations, we extracted a total of 5000 configurations out of the 60000 carbon structures generated in chapter 5. This data reduction was necessary at this stage to avoid over sampling of the C phase space during training. The C dataset was extracted in two steps: (i) we randomly select 12000 configurations from

the whole database (i.e 20 %) and perform k-mean clustering on the 12000 data with 50 centers (ii) sample a maximum number of configurations from each cluster such that on average, no one cluster center dominates the other in the final population.

6.4.3 First Principles Calculations

A single-point *ab initio* calculation is performed on all the structures visited during evolutionary algorithm configuration space search and MD refinement. The rVV10 [26] exchange and correlation functional was used and the same kinetic energy and charge density cutoffs as in chapter 5 are used. The Brillouin zone sampling was done with resolution of $0.014\times2\pi$ Å^{-1} for the all Li-C systems. All DFT calculations were performed with the Quantum ESPRESSO package [51, 52].

6.4.4 Atomic Environment Descriptors

The modified Behler and Parrinello descriptors [75, 97] as implemented in PANNA and described in chapter 4 are used to described atomic environment. We choose the same resolution of both radial and angular descriptors used to construct the general interatomic descriptors for C presented in chapter 5. Instead of the cutoff of 4.6 Å and 4.0 Å in chapter 5, we used a cutoff of 5.0 Å for both radial and angular part of the descriptors giving rise to 35 radial centers for the radial part and 18 radial centers and 8 angular centers for the angular part per species. Hence, the size of the descriptors becomes $35 \times 2 + 2 \times 3/2 \times 18 \times 8 = 502$. The cutoff distance of 5.0 Å was chosen to capture the Li-Li interaction on the surface of graphene sheet up to twice the lattice constant of graphene, the same as in the short-range part of the cluster expansion interatomic potential presented in chapter 3.

6.5 Results and Discussion

Here, we present our preliminary results on Li-C potential. This is the first iteration of the self-consistency cycle proposed in chapter 5 that comprises of a single step of evolutionary algorithm, clustering, molecular dynamics, *ab initio* calculations and neural network potential fitting. The training dataset is composed of 20000 Li-C data constructed by extracting, for each stoichiometry, configurations within 0.5 eV/atom from the lowest energy structures of that stoichiometry, 500 configurations of pure Li configuration with 0.05 eV/atom from the lowest energy configurations (because it is enough to predict the bulk Li in the body-centered cubic structure of Li) and finally 5000 pure C configurations sampled according to the description in Sec. 6.4.2. This gives a total of 25500 configurations out of which 20 % were set aside for validation and the remaining 80 % of the dataset was used for parameter optimization.

The parameters of the neural network models are optimized to reproduce *ab initio* total energy and the atomic forces.

6.5.1 Short Range Li-C Potential

We start with the analysis of the dependence of training accuracy on neural network architectures using the short range model. We examine different neural network architectures with fixed input vectors size of 502 and a single-node output layer as shown in Table 6.2. We used hyperbolic tangent as activation functions for the hidden layers and a linear function for the output layer.

Table 6.2: Training and validation with short range model as a function of increasing number of parameters. **Top**: energy root mean square error (RMSE) and mean absolute error (MAE). **Bottom**: root mean square error (RMSE) and mean absolute error (MAE) on forces per component. The columns denoted by Li, C and LiC represent errors evaluated on the test set corresponding to Li, C and LiC configurations separately, while column denoted by "All" is the error on all the validation dataset. Only an average training RMSE is reported. It is an average over the last 2500 optimization steps.

NN architectures	No. of parameters per species	training set (meV/atom)	test set (meV/atom)							
		RMSE	RMSE				MAE			
		All	Li	C	LiC	All	Li	C	LiC	All
502:64:32:1	34305	17 ± 3	2	35	21	21	2	24	13	13
502:64:32:16:1	34817	17 ± 3	2	31	19	19	2	21	12	12
502:128:64:32:1	**74753**	**13 ± 2**	**1**	**28**	**16**	**16**	**1**	**17**	**9**	**9**
502:256:64:32:1	147329	11 ± 3	1	26	15	15	1	16	7	7
502:256:128:64:32:1	172033	9 ± 2	1	28	16	16	1	16	8	8

NN architectures	No. of parameters per species	training set (meV/Å)	test set (meV/Å)							
		RMSE	RMSE				MAE			
		All	Li	C	LiC	All	Li	C	LiC	All
502:64:32:1	34305	204 ± 19	10	308	333	332	5	159	185	184
502:64:32:16:1	34817	210 ± 20	9	334	322	332	5	159	185	179
502:128:64:32:1	**74753**	**161 ± 17**	**6**	**411**	**316**	**315**	**3**	**149**	**151**	**150**
502:256:64:32:1	147329	139 ± 19	6	412	316	315	3	145	142	141
502:256:128:64:32:1	172033	119 ± 16	5	325	272	390	3	142	131	131

Table 6.2 shows the root mean square error (RMSE) and mean absolute error (MAE) on training and validation dataset. There is a noticeable dependence on the size of parameters which is more pronounced from the model with 34817 (502:64:32:16:1) to the model with 74753 (502:128:64:21:1) parameters. The accuracy of the models with 74753 parameters and greater are very similar and shows an overall improvement in accuracy over those with fewer parameters. This can further be confirmed from Fig. 6.2 where we show the correlation of a local quantity (change in force per component) between all the trained models and the model with 74753 parameters. From the figure, the reference model shows a better accuracy with respect to fewer parameters models

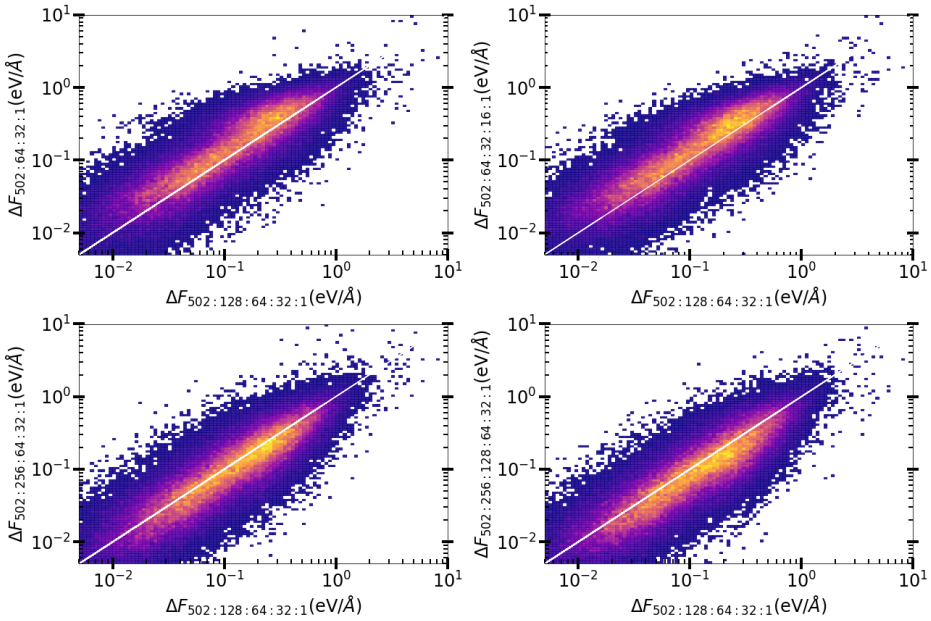

Figure 6.2: Representation of force deviation per component prediction on 2D histogram. All models are compared with the NNSR model. The axes labels represent the model under consideration.

(first row of Fig. 6.2) while a strong correlation can be observed between the reference model and those with more parameters (second row of Fig. 6.2). This improvement in accuracy can be seen both in energy and force predictions. Based on this analysis, the model architecture with 74753 parameters meets the balance between number of parameters and accuracy and it is therefore adopted for further analysis. We shall refer to it as NNSR in the rest of the discussion. Other short-range models will be explicitly specified whenever the need arises.

Overall, the energy RMSE in the training and validation dataset are of the same range, hence, there is no overfitting. For C configurations, the final validation RMSE of 35 meV/atom for the model 502:64:32:1 is slightly worst than the RMSE of 22 meV/atom reported in chapter 3 for C configurations. This may be related to the fact that some of the model parameters that were devoted to learn the C chemical environments alone are now adjusted to also learn the Li environments as well. Despite our expectations that long range electrostatics interaction may be needed to describe Li-C systems, surprisingly, an average energy RMSE of 14 meV/atom is observed, better than the average RMSE in C configurations that is expected to be well described by the short-range ANN methods.

In the case of the atomic forces, the overall validation RMSE in Table 6.2 is very similar to the force RMSE of 270 meV/Å observed in chapter 5. There is a noticeable difference between the training RMSE and the validation RMSE. In order to assess the source of this difference, we analyze the error distribution of force components on each component of the validation dataset independently (as done in Table 6.2). Fig. 6.3 shows the cumulative error distribution per force component. In all cases, more

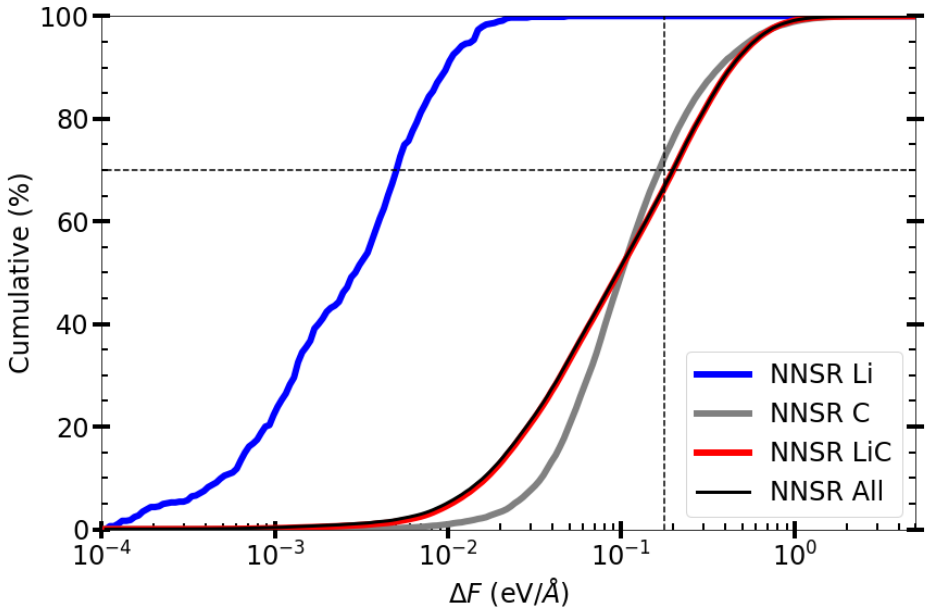

Figure 6.3: Cumulative error distributions on force components for the NNSR model: each point (x,y) on the curve signifies that y percent of the validation dataset have an error that is at most x. The vertical dashed line is the maximum training RMSE (i.e 161+17 meV/Å) while the horizontal dashed line represent 70 % cumulative.

than 70 % of the validation dataset have errors well below the training RMSE. Pure Li configurations are excellently predicted with more than 70 % of the configurations having error below 10 meV/Å. This is due to the fact that only very low energy configurations are included both in training and validation. The overall difference between the training and validation RMSE can be identified with the presence of outliers. The percentage of outliers may be reduced by continuing the self-consistency cycle proposed in chapter 5 (i.e, another step of, crystal structure search, clustering, molecular dynamics, *ab initio* calculations and training).

The excellent force prediction of Li-C configurations is even better than for C configurations, below 70 % cumulative combined with excellent prediction of energies of the Li-C configurations may lead to the conclusion that the long range electrostatics interaction is not required to describe Li-C systems. However, this conclusion is at variance with our experience from chapter 3 in which the cluster expansion potential with a pair interaction cutoff at 4.94 Å was not sufficient to fit the Li adsorption energy on graphene especially at low concentrations. A long-range electrostatic interaction correction was added for an accurate fit to be achieved. In order to assess the description of the long range effects in the NNSR model, we compare the prediction of adsorption energy of the Li on graphene substrate given by the NNSR model with the DFT values.

The Li adsorption energy is defined in equation 3.1 of chapter 3 with reference to the body-centered cubic phase of solid Li. In Table 6.3, we show the Li adsorption

Table 6.3: Li adsorption energy on graphene as a function of Li density: ΔE is the absolute deviation of the predicted energy from the DFT reference value computed at the same atomic coordinates. The prediction error on graphene is, $\Delta E_{\mathrm{C}}^{\mathrm{NNSR}} = 3.7\,\mathrm{meV/atom}$ and on Li is $\Delta E_{\mathrm{Li}}^{\mathrm{NNSR}} = 1.3\,\mathrm{meV/atom}$.

size	$d_{\text{Li-Li}}$(Å)	ΔE (meV/atom)		Binding Energy (eV/Li)	
		DFT	$\mathrm{NNP_{NNSR}}$	DFT	$\mathrm{NNP_{NNSR}}$
2x2	4.94	0.0	1.1	0.634	0.640
3x3	7.41	0.0	2.9	0.646	0.647
4x4	9.88	0.0	2.5	0.491	0.666
5x5	12.35	0.0	2.2	0.386	0.667
6x6	14.82	0.0	1.0	0.342	0.667
7x7	17.29	0.0	0.3	0.287	0.667
8x8	19.76	0.0	0.3	0.254	0.667

energy per Li as a function of increasing Li-Li distance (or decreasing Li concentration) as well as the prediction error per atom of the NNSR potential. Interestingly, the absolute deviation per atom denoted by ΔE on the table is less than 3 meV/atom for Li-C systems with the highest deviation being the pure C configurations with an error of 3.7 meV/atom. However, upon computing the Li adsorption energy, an excellent agreement can be seen for the 2x2 graphene supercell with a minimum Li-Li distance of 4.92 Å (smaller than the cutoff of 5.0 Å used for the environment descriptors). Also, the adsorption energy of Li in a 3x3 supercell, although having a minimum Li-Li distance larger than the cutoff, is still relatively well predicted. This is not surprising because the angular part of the descriptors permits interactions between atoms up to twice the cutoff. Indeed, at approximately twice the cutoff and beyond, the adsorption energy is wrongly predicted and remain constant within 1 meV/Li, an indication that the limit of the NNSR potential has been reached. Hence, confirming the need of a proper treatment of the long range interactions for accurate description of Li-C systems. In order to overcome the limitations of the NNSR model, we include the long range electrostatic treatments as described in Sec. 6.2.1 and implemented in the PANNA code. We present our preliminary results below.

6.5.2 Short Range with Long Range Li-C Potential

In order to include long-range electrostatics interactions as implemented in PANNA, we need to choose the Gaussian width of the atomic charge density and the species dependent atom hardness. Here, we choose Gaussian width as the covalent radius for Li and C (i.e, 1.28 and 0.76 Å for Li and C respectively). As for the atomic hardess, we choose 0.272 eV (or 0.1 Hartree) for Li and 0.544 eV (or 0.2 Hartree) for C.

Here, we choose a neural network architecture of 502:64:32:1 for both atomic energies and atomic electronegativity which gives rise to 68610 parameters per species which is of the same order as the number of parameters in the 502:128:64:32:1 NN

analyzed for the short range model. We choose the offsets of atomic electronegativities as 3.0 eV for Li and 6.3 eV for C corresponding to their atomic values reported in Ref. [160]. We train both the electronegativities and atomic energy concurrently with hyperbolic tangent activation functions for the hidden layers and a linear function for the output layer. This model, henceforth, will be referred to as NNLR1.

Table 6.4: Training and validation with short range with long range electrostatics correction. **Top**: energy root mean square error (RMSE) and mean absolute error (MAE). **Bottom**: root mean square error (RMSE) and mean absolute error (MAE) on forces per component. The columns denoted by Li, C and LiC represent errors evaluated on the test set corresponding to Li, C and LiC configuration separately, while column denoted by All is the error on all validation dataset. Only an average training error is reported. The average performed over the last 2500 optimization steps. The model names are written in parenthesis. NNSR and NNLR1 are already introduced as the short and long range models respectively. NNLR2 model as will be introduced later, corresponds to model similar to NNLR1 except that the output layer of the electronegativity is hyperbolic tangent function instead of a linear function.

NN architectures	No. of parameters per species	training set (meV/atom)	test set (meV/atom)							
		RMSE	RMSE				MAE			
		All	Li	C	LiC	All	Li	C	LiC	All
502:128:64:32:1 (NNSR)	74753	13 ± 2	1	28	15	14	1	17	8	8
502:64:32:1 (NNLR1)	68610	16 ± 3	2	30	20	20	2	22	12	12
502:64:32:1 (NNLR2)	68610	17 ± 3	1	32	19	19	1	20	11	11

NN architectures	No. of parameters per species	training set (meV/Å)	test set (meV/Å)							
		RMSE	RMSE				MAE			
		All	Li	C	LiC	All	Li	C	LiC	All
502:128:64:32:1 (NNSR)	74753	161 ± 17	6	411	296	295	3	149	136	135
502:64:32:1 (NNLR1)	68610	186 ± 22	7	409	329	329	4	160	174	174
502:64:32:1 (NNLR2)	68610	182 ± 25	5	342	315	314	3	157	173	172

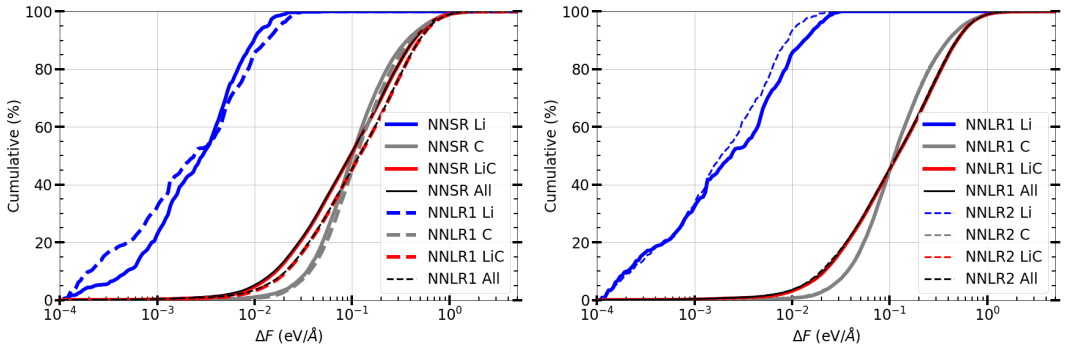

Figure 6.4: Cumulative error distributions on force components: each point (x,y) on the curve signifies that y percent of the validation dataset have an error that is at most x. **Left**: comparison of short range model with NNLR1 model: the dashed line legends designate the model trained with both long range and the solid lines correspond to short range model with 74753 parameters reported in Sec. 6.5.1. **Right**: comparison of NNLR1 and NNLR2 models: the solid lines are the curves for NNLR1 model and the dashed lines for the NNLR2 model.

In Table 6.4 we compare the accuracy of the NNLR1 model with the NNSR model.

The NNSR model highlighted in Table 6.2 and also reported in Table 6.4 is more accurate than the NNLR1 model. Similarly, the cumulative error distribution in Fig. 6.4 also confirms this. This superiority of the short range predictions may be due to difference in the total number of parameters per species between the NNSR model and the NNLR1 model. Indeed, a careful comparison of the accuracy of NNLR1 model with the model with architecture 502:64:32:1 with equivalent description for the short range atomic energies reveal a better agreement in training and validation RMSE.

Next, we examine the performance of the NNLR1 model on the Li adsorption energy on graphene surface as a function of supercell size. We present these results in Table 6.5 along side the results obtained with the NNSR model. Unlike the NNSR model that remains constant at large Li-Li distances, the NNLR1 captures the general trends of the adsorption energy despite having larger prediction error on these configurations. This is a further confirmation that the long range electrostatic interaction must be properly included in order to provide accurate description of the adsorption energy of Li in C materials.

Table 6.5: Li adsorption energy on graphene as a function of Li density: ΔE is the absolute deviation of the predicted energy from the DFT reference value computed at the same atomic coordinates. q_{Li} denotes the prediction of Li charge. The short range prediction error on graphene is, $\Delta E_{\mathrm{C}}^{\mathrm{SR}} = 3.7\,\mathrm{meV/atom}$ and on Li is $\Delta E_{\mathrm{Li}}^{\mathrm{SR}} = 1.3\,\mathrm{meV/atom}$. For the NNLR1 model, $\Delta E_{\mathrm{C}}^{\mathrm{NNLR1}} = 8.3\,\mathrm{meV/atom}$ and $\Delta E_{\mathrm{Li}}^{\mathrm{NNLR1}} = 3.4\,\mathrm{meV/atom}$.

size	$d_{\mathrm{Li\text{-}Li}}$(Å)	ΔE (meV/atom)			Binding Energy (eV/Li)			q_{Li}/e
		DFT	$\mathrm{NNP_{SR}}$	$\mathrm{NNP_{NNLR1}}$	DFT	$\mathrm{NNP_{SR}}$	$\mathrm{NNP_{NNLR1}}$	
2x2	4.92	0.0	1.1	4.2	0.619	0.640	0.644	0.137
3x3	7.41	0.0	2.9	5.3	0.635	0.647	0.681	0.213
4x4	9.88	0.0	2.5	3.0	0.464	0.666	0.627	0.237
5x5	12.35	0.0	2.2	3.2	0.368	0.667	0.618	0.246
6x6	14.82	0.0	1.0	4.2	0.325	0.667	0.615	0.249
7x7	17.29	0.0	0.3	4.8	0.274	0.667	0.613	0.251
8x8	19.76	0.0	0.3	5.3	0.237	0.667	0.612	0.252

In order to understand the origin of the slower decay of the Li adsorption energy predicted by NNLR1 with respect to DFT, we examine the charges on Li atom. As shown in the last column of Table 6.5, the general trend of the charges is similar to the one reported in Ref. [66] as well as our observations in chapter 3. However, the predicted charges are too small. For example, the charge on Li predicted by the NNLR1 model on a 4x4 supercell of graphene is 0.237 electrons which is more than twice smaller compared to a charge of about 0.54 electrons computed based on Hirshfeld charge decomposition techniques. It is also smaller than the 0.53 electrons based on Voronoi analysis reported in Ref. [66] or about 0.35 electrons with Löwdin analysis and about 0.9 electrons based on Bader analysis in Ref. [161]. Based on these evidences, we hypothesize that the slow decay of the adsorption energy is due to the underestimation of the Li charges by the NNLR1 model. To verify this hypothesis, we train another model with the same

number of parameters as NNLR1 model (referred to as NNLR2) where we replaced the linear activation function of the output layer for the electronegativity with a hyperbolic tangent. In this new model, the electronegativity prediction can only deviate from isolated atom value by ± 1 since the offset of the electronegativities are set to the atomic values. In what follows, we compare the performance of the NNLR1 and NNLR2 models on the validation dataset and on the adsorption energy.

The training and validation error in energy and forces of the NNLR1 and NNLR2 models are very similar as shown in Table 6.4. This observations can also be seen from Fig. 6.4 in which the prediction in C and Li-C systems are basically identical while a slight difference in the prediction of Li configurations can be noticed. Considering that the charge on Li in the two models for Li-C systems are completely different (see Table 6.6) and yet a similar energy and force prediction error are achieved in the validation set is an indication that the long range contribution in Li-C systems may be too small to be reliably identified from the global RMSE minimization and could be completely obscured by the short-range model error as seen in Sec. 6.5.1.

Table 6.6: Li adsorption energy on graphene as a function of Li density: ΔE is the absolute deviation of the predicted energy from the DFT reference value computed at the same atomic coordinates. q_{Li} denotes the prediction of Li charge. The prediction error for the NNLR1 model on graphene and Li-bulk are respectively, $\Delta E_{\mathrm{C}}^{\mathrm{NNLR1}} = 8.3\,\mathrm{meV/atom}$ and $\Delta E_{\mathrm{Li}}^{\mathrm{NNLR1}} = 3.4\,\mathrm{meV/atom}$. For NNLR2, they are $\Delta E_{\mathrm{C}}^{\mathrm{NNLR1}} = 4.6\,\mathrm{meV/atom}$ and $\Delta E_{\mathrm{Li}}^{\mathrm{NNLR1}} = 1.0\,\mathrm{meV/atom}$.

size	$d_{\text{Li-Li}}$(Å)	ΔE (meV/atom)			Binding Energy (eV/Li)			q_{Li}/e	
		DFT	$\mathrm{NNP_{NNLR1}}$	$\mathrm{NNP_{NNLR2}}$	DFT	$\mathrm{NNP_{NNLR1}}$	$\mathrm{NNP_{NNLR2}}$	$\mathrm{NNP_{NNLR1}}$	$\mathrm{NNP_{NNLR2}}$
2x2	4.92	0.0	4.2	2.7	0.619	0.644	0.608	0.137	0.336
3x3	7.41	0.0	5.3	5.6	0.635	0.681	0.447	0.213	0.532
4x4	9.88	0.0	3.0	1.7	0.464	0.627	0.377	0.237	0.595
5x5	12.35	0.0	3.2	4.0	0.368	0.618	0.343	0.246	0.615
6x6	14.82	0.0	4.2	4.6	0.325	0.615	0.330	0.249	0.623
7x7	17.29	0.0	4.8	5.0	0.274	0.613	0.322	0.251	0.627
8x8	19.76	0.0	5.3	5.2	0.237	0.612	0.319	0.252	0.629

To measure the correlation between charge prediction and the accuracy of the resulting adsorption energy prediction, we investigate the performance of NNLR1 and NNLR2 models on Li adsorption energy on graphene substrate. Table 6.6 shows the adsorption energy, the prediction error and the predicted charges of the NNLR1 model and the NNLR2 model. With the constraint imposed on the electronegativity in NNLR2 model, we observed an improvement on the charge prediction towards the Hirshfeld charge of 0.54 electrons for Li adsorption on 4x4 graphene supercell. The trend in the adsorption energy is maintained (except from 2x2 to 3x3 supercell) and a better prediction of the adsorption energy is attained especially at low concentrations (large Li-Li distances). Remarkably, we see a systematic improvement in the description of the adsorption energy from NNSR potential with charge, $q_{Li} = 0$, to NNLR1 potential with charge, $q_{\mathrm{Li}}^{\mathrm{NNLR1}}$, to NNLR2 with charge $q_{\mathrm{Li}}^{\mathrm{NNLR2}} \approx 3 \times q_{\mathrm{Li}}^{\mathrm{NNLR1}}$, showing the dependence of the long range description on the accurate prediction of the electronegativities

and hence of the atomic charges.

The results obtained thus far lead to the conclusion that an accurate description of Li adsorption energy in graphene (and in general, a generic C environments) is not based on accurate fit alone but a combination of accurate fit and proper decomposition of the total energy into long range and short range contributions. This is particularly crucial in Li-C systems since the long range contribution is smaller than a typical prediction error of the artificial neural network potentials of a few meV/atom and therefore, the long range contribution is basically treated as noise during the optimization process. In order to circumvent this, a proper amplification of this component of the total energy is needed. This amplification can be achieved by choosing a charge density decomposition scheme such as the Hirshfeld method [35] and constrain the model parameters also on the atomic charges in addition to the total energy and atomic forces. We believe that this constraints on atomic charges may not be necessary to accurately describe systems with more significant long range electrostatic contributions such as molten salt or polar molecules.

Chapter 7

Conclusions and Outlooks

In the course of this thesis, we have made attempts to elucidate the interactions of Li with carbon materials that may be responsible for the experimentally observed high lithium density on carbon materials by developing inexpensive and highly accurate interatomic potentials to describe such systems. We started with the site-based (cluster expansion) interatomic potential method that is suitable to describe interactions in well defined lattice systems and then proceeded to the artificial neural network techniques that provides a framework to construct an interatomic potentials for generic systems in its currently used form and then its extension to include long range electrostatic interactions.

We presented a detailed study of the interaction of Li with single layer graphene both at zero temperature within Density Functional Theory (DFT) and at finite temperature via a combination of cluster expansion techniques and Grand Canonical Monte Carlo methods. The energetic stability of Li adatoms (isolated Li atoms) and Li-clusters on graphene substrate was studied. While at zero temperature, both the Li adatoms and small Li-clusters were found to be unstable with respect to phase separation into graphene and Li-bulk, thermal effects were found to be sufficient to prevent phase separation of the Li-graphene system at low concentrations at finite temperatures. However, above the stoichiometry of $\approx LiC_6$, we found formation of Li-clusters to be energetically favorable with respect to random distribution of Li atoms. We predicted two stable phases of two dimensional (2D) Li adsorption on graphene substrate (Li-gas and Li-cluster) and the range of concentration (i.e. the chemical potential) at which the two phases are in equilibrium. This coexistence region was found to span about one-third of the entire concentration range at low temperatures. Within the Li-cluster phase, a phase where islands of Li join and form continuous stripes was also observed which however, by varying the size of the simulation cell, we showed to be due to finite size effects and would eventually vanish in the thermodynamic limit. Although a simulation artifact in the case of an infinite graphene sheet, we speculated that such phase can occur at the edges of a graphene flake and seeds the lithiation.

Due to the limitations of the cluster expansion method, we adopted the artificial neural network (ANN) technique that is suitable to construct general purpose interatomic potentials within a machine learning approach. In order to be able to construct the general interatomic potential for Li-C systems, we implemented an all-to-all fully connected feed-forward ANN machine learning methods based on Behler and Parrinello symmetry functions and their modifications in the "Properties from Artificial Neural Network Architectures" (PANNA) package. We tested its features and capabilities by constructing an interatomic potential for organic molecules with an ANN architecture with about half a million parameters, the same architecture used in Ref. [97], based on large dataset reported in Ref. [102]. Notably, we reproduced the results [97] excellently well thereby validating the capability of the package.

As the flexibility and transferability of ANN interatomic potentials depend strongly on the dataset used in their training, we proposed an automated, system-independent, self-consistent technique that integrates evolutionary algorithm (EA), unsupervised machine learning in the form of clustering, and molecular dynamics (MD) to generate a diverse dataset suitable to construct a truly general interatomic potential. We demonstrated its successful application to the challenging case of carbon for which classical and machine-learnt potentials are abundant in literature and showed that MD simulation combined with EA and clustering successfully explore a wide range of configurations on equal footing so that the resulting dataset covers energy and volume landscape rather homogeneously. The final neural network potential (NNP) was found to predicts structural and elastic properties of diamond and graphene very accurately and also ranked accurately, a diverse set of sp^3 C configurations. When compared to classical force-fields such as Tersoff and ReaxFF, our NNP out-performed them and more interestingly, was found to be better than the existing machine learnt potentials on average. Furthermore, our NNP predicts the majority of the graphite phonon modes excellently although it predicts instability (imaginary frequencies) for the very soft modes that are related to interlayer interaction. We have traced this behaviour to the accuracy requirement in predicting such small forces. We demonstrated two possible ways to alleviate this problems. The first approach involves training the NNP with a relative force deviation that weights more small atomic forces during parameter optimization while the second approach was to build the training set only with structures that are in the vicinity of graphite according to a fingerprint-based distance measure. Remarkably, the first approach yielded accurate frequencies for the soft phonon modes with slightly worst prediction of the high frequency modes. However, a further tuning of training parameters and the loss function can be used to construct a high accuracy model. The second approach was found to provide accurate phonon frequencies for the soft modes but showed poor generalization to a wider range of structures, compared to a more comprehensive potential trained on the entire dataset.

Finally, we presented our extension of the PANNA package to include the long range electrostatic interaction through the charge equilibration scheme in which a network learns environment-dependent atomic electronegativities in addition to the local atomic energies. We presented preliminary results of a general interatomic potential for Li-C systems using data generated based on our proposed self-consistent approach. We constructed Li-C interatomic potential with and without long range treatments. The short range Li-C interatomic potential predicted the total energy and forces relatively accurately but failed to described the general trend of the adsorption energy of Li adsorbed on graphene substrate. Instead, including the long range electrostatics although less accurate compared to the short-range model, captures the general trend of the Li-adsorption energy on graphene substrate but decays slowly compared to DFT values. The lower average accuracy of the long-range model compared to the short range model was ascribed to the differences in the number of parameters which is higher in the short-range model than in the long-range model. The slow decay of the Li-adsorption energy was found to be due to too small charge predictions by the long range potential in comparison to the DFT values. We verified this by constructing a second neural network in which we constrained the atomic electronegativity within ± 1 from its atomic value thereby leading to larger charges. The second long-range potential has a similar accuracy as the first one on the validation dataset but predicts both the Li-adsorption energy as well as the atomic charges more accurately. The results lead us to conclude that the long-range electrostatic interaction in Li-C systems is comparable to the typical overall prediction error attained by the machine learning based interatomic potential and therefore requires a special amplification strategy to extract the long-range interaction from the total energy, so as to enhance the signal to noise ratio for the long range component. The need for the development of an amplification strategy suitable to accurately extract the long-range components from total energy necessitates further investigations.

Currently, we are investigating the dependence of the long-range ANN model on the number of parameters. This investigation is aimed to shed more light on the dependence of the signal to noise ratio of the long-range component on the flexibility of the ANN model.

In the short-term, we plan to extend the current long-range PANNA package to constrain the ANN model parameters so as to reproduce DFT atomic charges and train a ANN potential for the Li-C systems to reproduce the Hirshfeld charges in addition to the current constraints on the energy and the atomic forces. This decomposition is needed to correctly extract the electrostatic interactions at long distances. Concurrently, we will construct an ANN potential for the organic molecules used in chapter 4 with our current long range PANNA package and compare results with the one reported in chapter 4.

Once an accurate Li-C potential is constructed, atomistic simulations can be performed to provide hints on several unanswered questions about the high Li density in graphitic anode materials. In the long-term, we plan to resolve the following open problems: (i) the combined effects of defects, edges and vacancies in the stability of high Li-density in graphitic materials, (ii) thermal effects on Li adsorption and intercalation in layered C systems, (iii) the combined effects of defects, edges, vacancies and thermal effects on very large C systems at/or closed to sizes of graphitic flakes observed in literatures and so on.

In the current model, the explicit dependence of the short range part of equation 6.1 on atomic charges is neglected. We plan to extend our model to include this dependence.

Appendix A

Lithium Adsorption on Graphene at Finite Temperature

A.1 Orthogonality Condition of Basis Functions

The cluster expansion technique is rooted on the fact that a complete basis functions of occupation, $\Phi_f(C)$ exists such that any function of occupation can be written as a linear combination of those basis functions. These basis functions are constructed from the product of occupation number of site $i = 1, ..., N$ of an N sites lattice which reads as

$$\Phi_f(C) = \prod_{i \in f} c_i, \tag{A.1}$$

where $C = \{c_i\}, i = 1, ..., N$. The basis functions satisfied the following orthogonality conditions:

$$\sum_C \Phi_f(C)\Phi_{f'}(C) = 2^{-N}\delta_{ff'}, \tag{A.2}$$

and implies

$$\sum_f \Phi_f(C)\Phi_f(C') = 2^{-N}\delta_{CC'}. \tag{A.3}$$

In this work, we choose and equivalent definition where the occupation number of site i is 1 if the site is occupied by Li and 0 if the site is empty. The total energy of a given configuration C is then given by;

$$E_{CE}(C) = \sum_f J_f \Phi_f(C) \tag{A.4}$$

where linear coefficients, J_f are called the effective cluster interactions (ECI). There is one J_f for each cluster (also called figure) which are to be determined via linear regression. We refer to *cluster* or *figure* as a collection of lattice sites.

A.2 Monitoring the Fitting Procedure

Fitting is performed by adding a configuration with the worst prediction error to the training set. This means that at each step, a configuration with the most information needed to improve the fit is added. Indeed after few steps, the root-mean square error converge and subsequent configurations do not provide information that can significantly modify the parameters of the model.

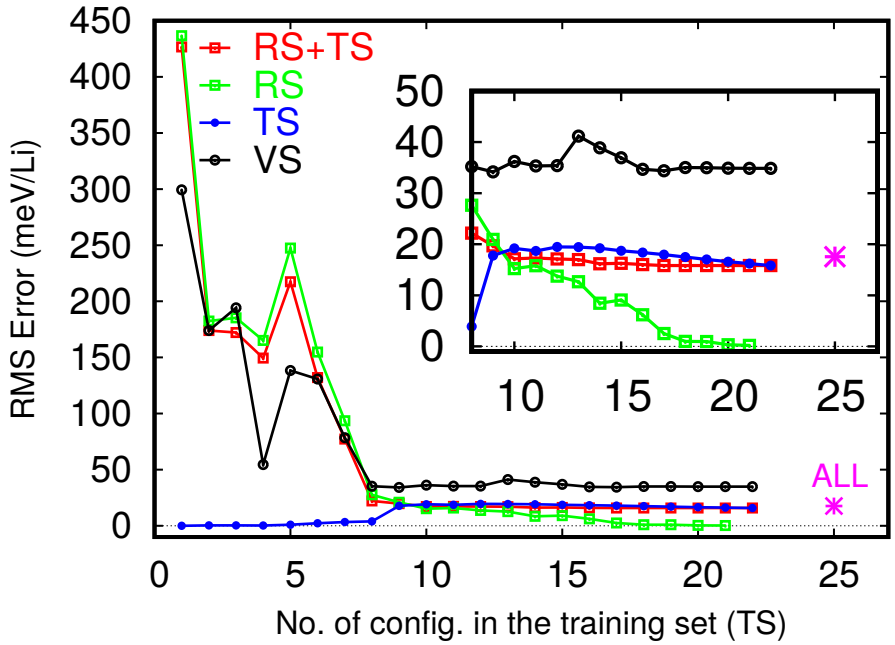

Figure A.1: Evolution of the root-mean-square (RMS) error as the number of configurations used in fitting, i.e. in the training set (TS), increases. In this analysis, we use a total of 25 different configurations of Li on graphene. Three configurations are set aside as validation set (VS) and are never included during the training of the model. Therefore, at zeroth step of the training, we have zero configurations in the training set (TS) and 22 configurations in the remaining set (RS). At each step of the training the configuration with the largest prediction error is added to TS, reducing the number of configurations in the RS by one. We continue these training steps until all 22 configurations in RS have been added to TS. Finally, for analysis purposes we also perform a training including all 25 configurations (magenta asterisk labeled as ALL). The inset shows the zoom of the plot at later steps of training where RMS error converges for TS.

A.3 Performance of the Model

In Fig. A.2 the deviation of the DFT energy from the energy predicted by the full model described in the manuscript is shown. The root-mean-square deviation is 18.0 meV/Li (1.6 meV/C) with a cross validation error of 30.2 meV/Li (3.2 meV/C) as already reported in the main text.

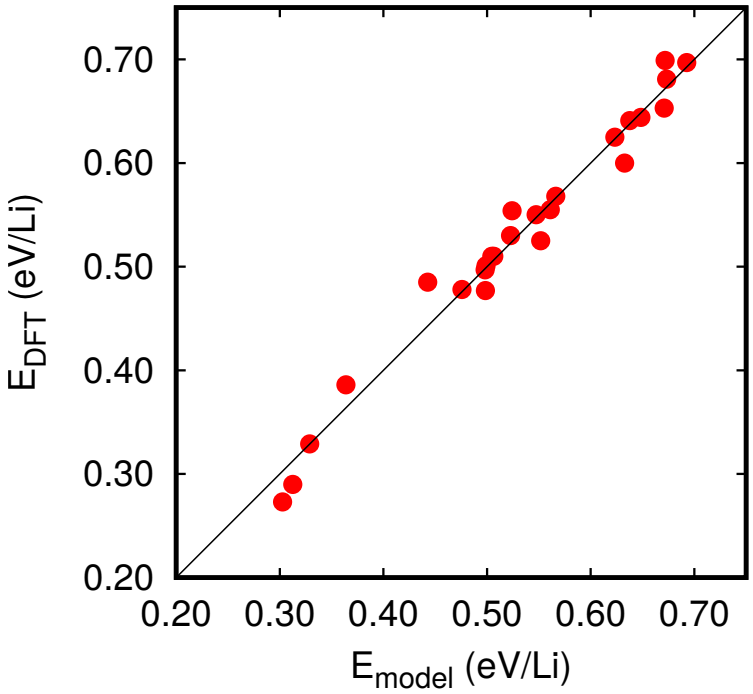

Figure A.2: Comparison of the model energies with DFT ones using rVV10 functional. The black line represents perfect agreement while the red points are the actual model predictions.

Table A.1: Best fit values for the effective cluster interactions of the model when 3- and 4-body interaction terms are neglected. All parameters, except γ, in eV.

	J_1	J_{2a}	J_{2b}	J_{2c}	J_3	J_4	J_{dd}	γ
ECI	0.429	0.074	-0.177	0.051	0.000	0.000	1.198	0.559

When 3- and 4-body terms are neglected, the obtained results are similar to the one reported in Ref. 8: the most influential cluster figure is the point cluster as shown in table A.1. The resulting root-mean-square deviation is 98.3 meV/Li (7.5 meV/C) with a cross validation error of 158.2 meV/Li (21.2 meV/C). The reduced quality of the model is shown in Fig. A.3.

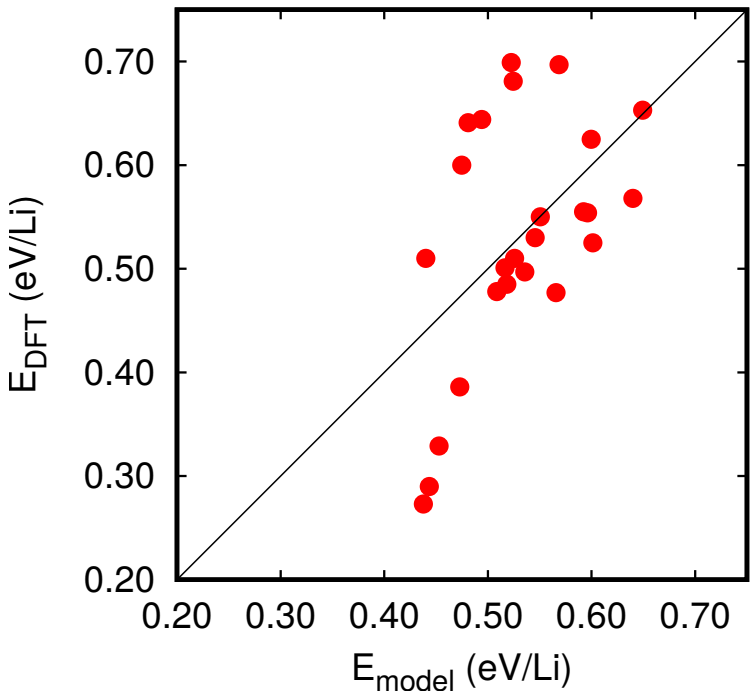

Figure A.3: Comparison of the model energies with DFT ones using rVV10 functional when 3- and 4-body interaction terms are neglected. The black line represents perfect agreement while the red points are the actual model predictions.

Even less accurate results are obtained when in addition, the J_{dd} term is set to zero. Results are reported in table A.2. The root-mean-square deviation is 108.3 meV/Li (7.3 meV/C) with a cross validation error of 133.9 meV/Li (15.3 meV/C). The DFT *vs* model comparison is shown in Fig. A.4. This is a strong indication that the long-range dipole-dipole term and the 3- and 4- body interactions are very important for a proper description of Li adsorption on graphene.

Table A.2: Best fit values for the effective cluster interactions of the model when 3- and 4-body and long-range interaction terms are neglected. All parameters, except γ, in eV

	J_1	J_{2a}	J_{2b}	J_{2c}	J_3	J_4	J_{dd}	γ
ECI	0.486	0.068	-0.077	0.045	0.000	0.000	0.0000	0.000

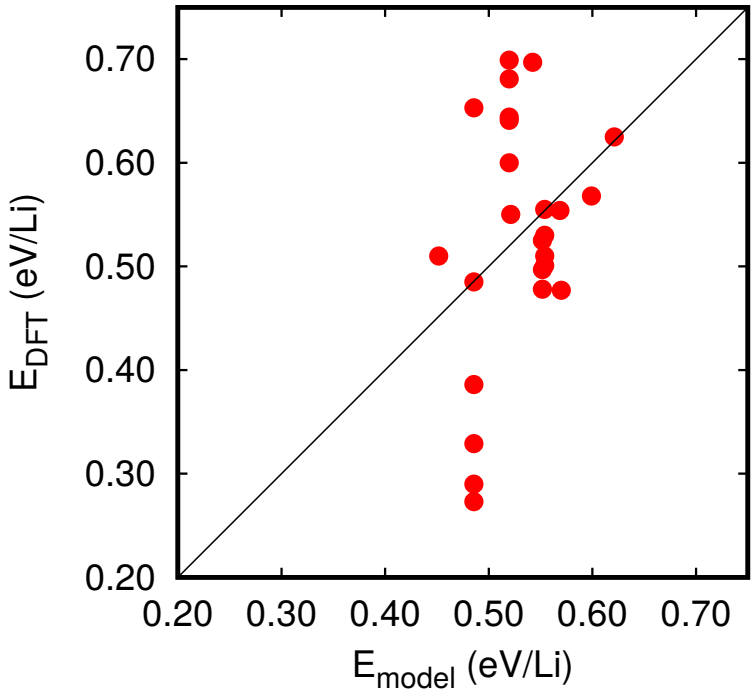

Figure A.4: Comparison of the model energies with DFT ones using rVV10 functional when 3- and 4-body and long-range interaction terms are neglected. The black line represents perfect agreement while the red points are the actual model predictions.

When only the dipole-dipole long-range terms are neglected the resulting parameters are the ones reported in table A.3 (DFT *vs* model deviations shown in Fig. A.5). This model describes relatively well the Li-dense configurations, but dilute configurations with Li adatoms in 3x3 to 8x8 are poorly described indicating the importance of the long-range dipole-dipole interaction term.

The root-mean-square deviation is 70.9 meV/Li (5.2 meV/C) with a cross validation error of 91.9 meV/Li (8.6 meV/C)

Table A.3: Best fit values for the effective cluster interactions of the model when only the dipole-dipole long-range terms are neglected. All parameters, except γ, in eV

	J_1	J_{2a}	J_{2b}	J_{2c}	J_3	J_4	J_{dd}	γ
ECI	0.408	0.491	-0.037	0.065	-1.130	0.196	0.0000	0.000

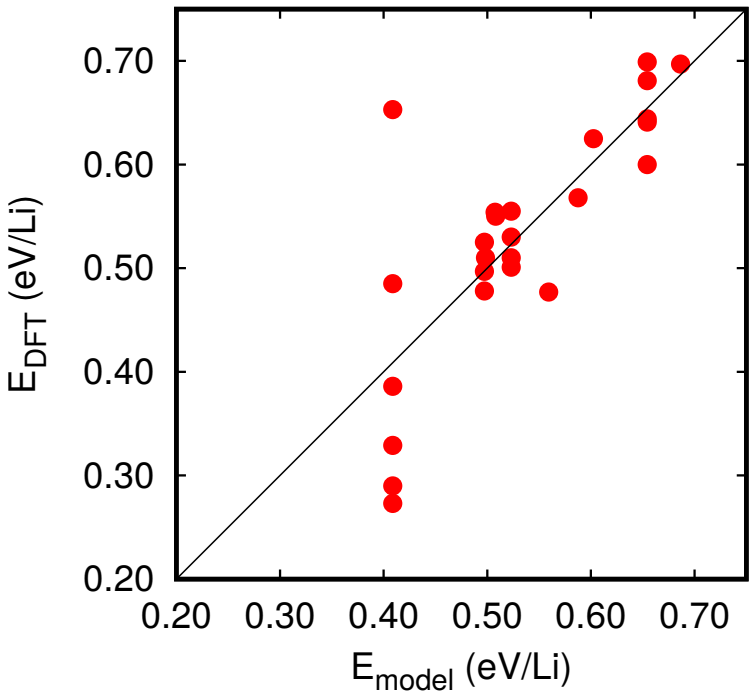

Figure A.5: Comparison of the model energies with DFT ones using rVV10 functional when only the long-range interaction terms are neglected. The black line represents perfect agreement while the red points are the actual model predictions.

A.4 Cluster Expansion : The behaviour of the dipole-dipole interaction term

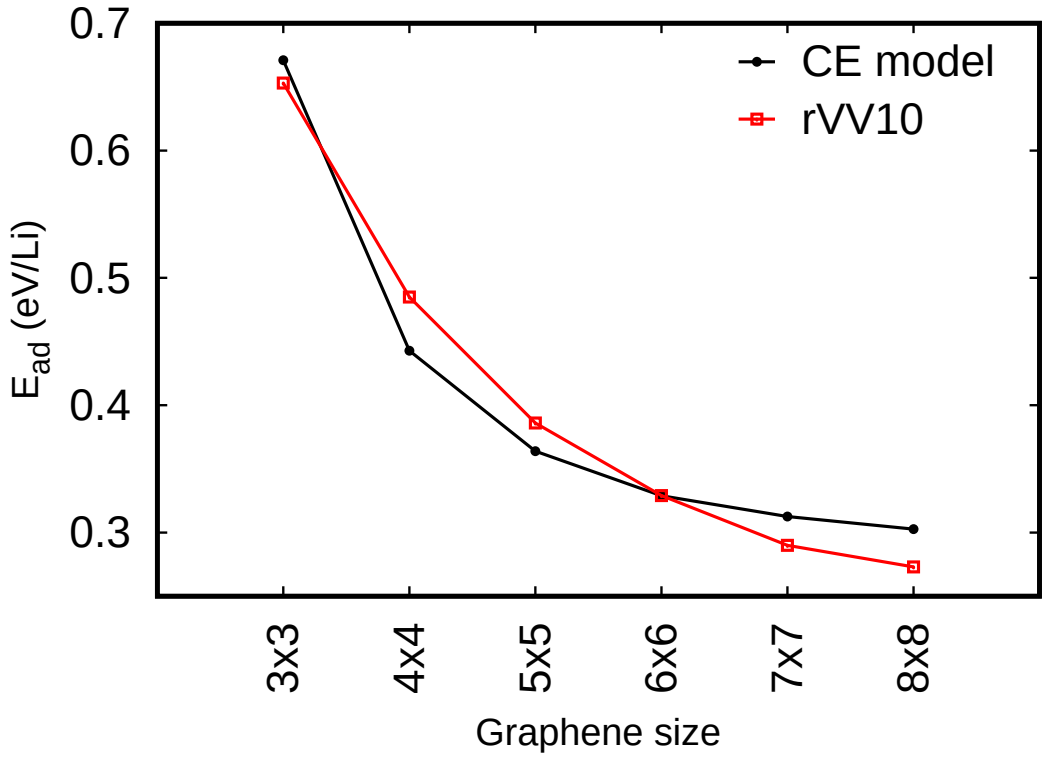

Figure A.6: The behaviour of the dipole-dipole part of the cluster expansion model obtained via single Li adatom on various supercell sizes. Note that in all the supercell sizes used, Li ions in adjacent cells are far enough that the only contributions to the model energy are the onsite term, which constitutes a constant shift, and the long-range interaction term. The model can be seen to capture the long-range behaviour observed in the DFT calculations (red squares)

A.5 Thermodynamics Integration in Details

For a given inverse temperature β and the chemical potential μ, the grand canonical thermodynamics potential per site is given by

$$\Phi(\beta,\mu) = -\frac{1}{\beta N_s} log\ Q(\beta,\mu) \tag{A.5}$$

Where N_s is the number of sites and $Q(\beta,\mu)$ is the grand canonical partition function given by

$$Q(\beta,\mu) = \sum_{N=0}^{N_s} \sum_{\{c\}} \exp(-\beta[E(C) - \mu N]) \tag{A.6}$$

The Helmholtz free energy, $\mathrm{F}(\beta, x)$ per site is the Legendre transform of $\Phi(\beta,\mu)$ and written as $F = \Phi + \mu x$.

The total differential of equation A.5 is given by

$$d(\beta\Phi) = (\frac{E}{N_s} - \mu x)d\beta - \beta x d\mu \tag{A.7}$$

The integration of equation A.7 can be performed along a continuous path in the (β,μ) space. Thus for a fixed β path, Φ is

$$\Phi(\beta_0, \mu) = \Phi(\beta_0, \mu_0) - \int_{\mu_0}^{\mu} d\mu x(\beta, \mu) \tag{A.8}$$

and along the constant μ path we have,

$$\beta\Phi(\beta, \mu_0) = \beta_0\Phi(\beta_0, \mu_0) + \int_{\beta_0}^{\beta} d\beta(\frac{\langle E\rangle}{N_s} - \mu x) \tag{A.9}$$

The values of $\langle E\rangle$ and x are easily obtain from a grand canonical Monte Carlo simulation at any given values of β and μ. Once obtained, the integration of equations A.8 and A.9 can be performed with any integration method of choice. In this work, integration along the constant temperature path was performed with trapezoidal integration scheme. The integrand along the constant μ part contains two fluctuating variables. To minimize the error in the integration, we approximate the integrand with a polynomial of degree 3. The Monte Carlo data were fitted to the polynomial to obtain the coefficients and then perform the integration over the inverse temperature β. We checked that polynomial of degree 3 was sufficient for the fit by comparing the result of the fit with polynomial of degree 2 and 4.

In the constant μ path, we write the integrand in equation A.9 as a power series in the inverse β. The integrand is given by

$$E - \mu_0 x = b_1 + \frac{b_2}{\beta} + \frac{b_3}{\beta^2} + \frac{b_4}{\beta^3} \tag{A.10}$$

where b_i are functions of μ_0. Equation A.9 becomes

$$\Phi(\beta, \mu_0) = \frac{\beta_0}{\beta}\Phi(\beta_0, \mu_0) + b_1(\beta - \beta_0) + \frac{b_2}{\beta} log\left(\frac{\beta}{\beta_0}\right) - \frac{b_3}{\beta}(\frac{1}{\beta} - \frac{1}{\beta_0}) - \frac{b_4}{2\beta}(\frac{1}{\beta^2} - \frac{1}{\beta_0^2}) \tag{A.11}$$

To compute the energy difference between two phases, we must integrate equation A.7 along a continuous path in the (μ, β) space. The energy difference between two points in the (μ, β) space characterized by (μ_1, β_1) and (μ_2, β_2) respectively is given along continuous path as follows

$$\beta_2\Phi(\beta_2, \mu_2) - \beta_1\Phi(\beta_1, \mu_1) = \int d(\beta, \mu)\left(\frac{\langle E\rangle}{N_s} - \mu x, -\beta x\right) \tag{A.12}$$

Integrating equation A.12 gives

$$\begin{aligned}\beta_2\Phi(\beta_2,\mu_2)-\beta_1\Phi(\beta_1,\mu_1) &= \int_{\beta_1}^{\beta_0} d\beta\left(\frac{\langle E\rangle}{N_s}-\mu_1 x\right)+\int_{\mu_1}^{\mu_2}(-\beta_0)x d\mu \\ &+ \int_{\beta_0}^{\beta_2} d\beta(\frac{\langle E\rangle}{N_s}-\mu_2 x) \end{aligned} \tag{A.13}$$

At constant temperature, i.e $\beta_1=\beta_2$, equation A.13 becomes

$$\begin{aligned}\beta\Delta\Phi(\beta) &= \int_{\beta}^{\beta_0} d\beta(\frac{\langle E\rangle}{N_s}-\mu_1 x)+\int_{\mu_1}^{\mu_2}(-\beta_0)x d\mu \\ &+ \int_{\beta_0}^{\beta} d\beta(\frac{\langle E\rangle}{N_s}-\mu_2 x) \end{aligned} \tag{A.14}$$

where $\Delta\Phi=\Phi(\beta,\mu_2)-\Phi(\beta,\mu_1)$ is the free energy difference at temperature β between two phases defined by the chemical potential μ_1 and μ_2. Using equation A.10, we have

$$\begin{aligned}\Delta\Phi &= (b_1^{\mu_2}-b_1^{\mu_1})(1-\frac{T}{T_0})+(b_2^{\mu_2}-b_2^{\mu_1})k_BTlog(\frac{T_0}{T}) \\ &- (b_3^{\mu_2}-b_3^{\mu_1})k_BT(k_BT-k_BT_0)-\frac{1}{2}(b_4^{\mu_2}-b_4^{\mu_1})k_BT((k_BT)^2-(k_BT_0)^2) \\ &- \frac{T}{T_0}\int_{\mu_1}^{\mu_2} x(T_0)d\mu \end{aligned} \tag{A.15}$$

Thus, the change in Helmholtz free energy per site at a constant temperature between the concentrations $x(\mu_1)$ and $x(\mu_2)$ is given by

$$\Delta F(\beta)=\Delta\Phi+\mu_2 x(\mu_2)-\mu_1 x(\mu_1) \tag{A.16}$$

A.6 Error Analysis

To estimate the error in free energy, let's the expand the free energy around β and μ for a given change $d\beta$ and $d\mu$ to second order, for convenience, we expand $\beta\Phi(\mu+d\mu,\beta+d\beta)$ so that

$$\begin{aligned} d(\beta\Phi) &= \frac{\partial(\beta\Phi)}{\partial\beta}d\beta+\frac{\partial(\beta\Phi)}{\partial\mu}d\mu \\ &+ \frac{1}{2}[\frac{\partial^2(\beta\Phi)}{\partial\beta^2}(d\beta)^2+\frac{\partial^2(\beta\Phi)}{\partial\mu^2}(d\mu)^2+2\frac{\partial}{\partial\beta}\frac{\partial}{\partial\mu}(\beta\Phi)d\mu d\beta] \qquad \text{(A.17)} \\ &= (\frac{E}{N_s}-\mu x)d\beta-\beta x d\mu+d(\beta\epsilon_\Phi) \qquad \text{(A.18)} \end{aligned}$$

Comparing equation A.7 and A.17 we see that

$$d(\beta\epsilon_\Phi) = \frac{1}{2}\left[\frac{\partial^2(\beta\Phi)}{\partial\beta^2}(d\beta)^2 + \frac{\partial^2(\beta\Phi)}{\partial\mu^2}(d\mu)^2 + 2\frac{\partial}{\partial\beta}\frac{\partial}{\partial\mu}(\beta\Phi)d\mu d\beta\right] \tag{A.19}$$

is the error in $d(\beta\Phi)$. Thus the error in Φ becomes

$$\epsilon_\Phi = \frac{1}{\beta}\int\frac{1}{2}\left[\frac{\partial^2(\beta\Phi)}{\partial\beta^2}(d\beta)^2 + \frac{\partial^2(\beta\Phi)}{\partial\mu^2}(d\mu)^2 + 2\frac{\partial}{\partial\beta}\frac{\partial}{\partial\mu}(\beta\Phi)d\mu d\beta\right] \tag{A.20}$$

Along constant temperature path, we have

$$\epsilon_\Phi = \frac{1}{2}N_s\beta|d\mu|\int\epsilon_x^2 d\mu \tag{A.21}$$

And along the constant μ path,

$$\epsilon_\Phi = -\frac{1}{2}N_s\frac{|d\beta|}{\beta}\int\left[\epsilon_e^2 - \mu^2\epsilon_x^2\right]d\beta \tag{A.22}$$

Where $\epsilon_x^2 = \langle x^2\rangle - \langle x\rangle^2$ and $\epsilon_e^2 = \langle\frac{E^2}{N_s^2}\rangle - \langle\frac{E}{N_s}\rangle^2$ are the variance of x and $\frac{E}{N_s}$ respectively

A.7 Phases Identification

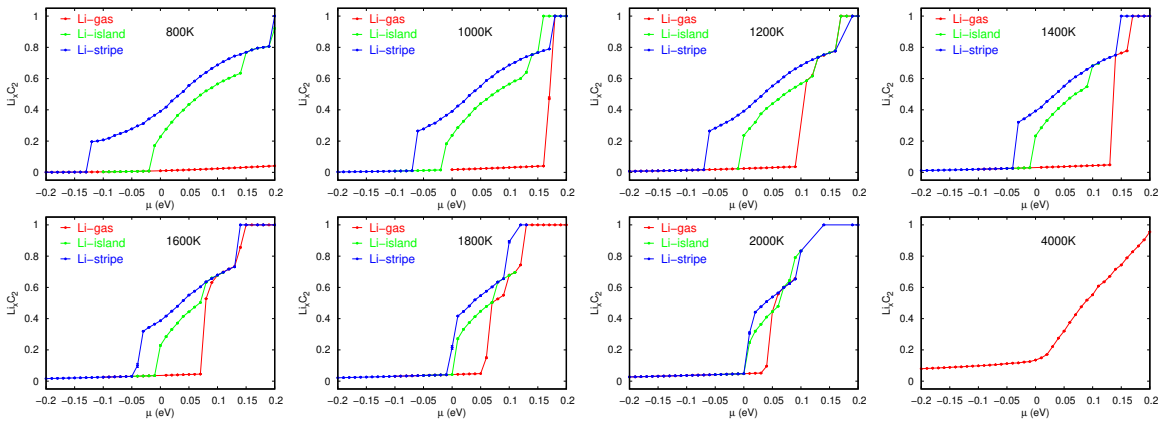

Figure A.7: Concentration,x as a function of chemical potential, μ: The color codes indicate the phases. As we scan the chemical potential, we go from one phase (x is continues) to another phase identified with a jump in the x. Here, we show the results from 800 K to 2000 K at a step of 200 K and include also the plot at 4000 K reference where a continues transition between phases can be observed.

Appendix B

Appendix

B.1 Self consistent exploration of the phase space

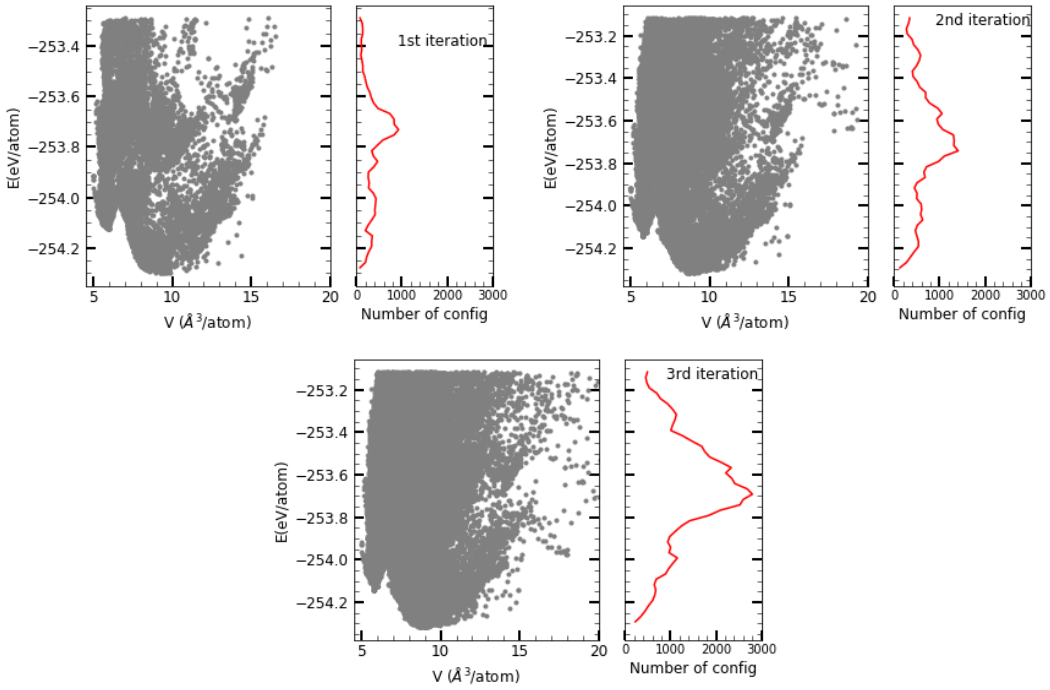

Figure B.1: The per atom energy-volume distribution of sampled configurations at each self consistency iteration (left column); and the corresponding number analysis with respect to the energy range sampled (right column).

B.2 Relative loss versus standard loss

Table B.1: Comparison of elastic constants of diamond trained with the quadratic force loss (NNP) and relative force loss (NNP retrain). See Table II in main text for further details on the method of calculation.

DIAMOND	NNP	NNP retrain	DFT
a (Å)	3.576	3.576	3.584
B_0 (GPa)	431	511	425
C_{11} (GPa)	1054	1106	1044
C_{12} (GPa)	119	213	116
C_{44} (GPa)	542	513	547
GRAPHITE	NNP	NNP retrain	DFT
a (Å)	2.471	2.466	2.471
c (Å)	6.732	6.681	6.719
B_0 (GPa)	48	88	40
C_{11} (GPa)	1053	1016	1048
C_{12} (GPa)	197	210	182
C_{13} (GPa)	-23	-2	-5
C_{33} (GPa)	57	103	43
C_{44} (GPa)	-5	4	4
C_{66} (GPa)	428	403	433
GRAPHENE	NNP	NNP retrain	DFT
a (Å)	2.470	2.471	2.470
ν	0.244	0.209	0.173
E (GPa)	967	923	1015
C_{11}(GPa)	1028	965	1047
C_{12} (GPa)	251	202	181

B.3 Amorphous carbon

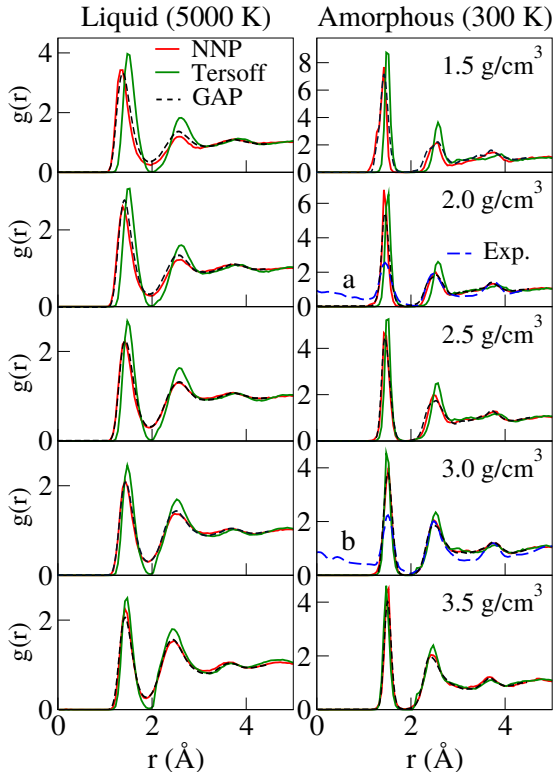

Figure B.2: Radial distribution function for liquid and amorphous carbon, equivalent to Fig.6a of the main text, with the addition of results for the GAP potential (digitized from Ref. [91]) and experimental results from Ref. [162] (blue dashed line-a in the figure, at 2.0 g/cm^3) and Ref. [163] (blue dashed line-b in the figure, at 3.0 g/cm^3). Experimental curves were scaled and vertically shifted to match the large r behaviour, and processed to obtained $g(r)$ as needed in Ref. [162]

B.4 Amorphous and Liquid Carbon

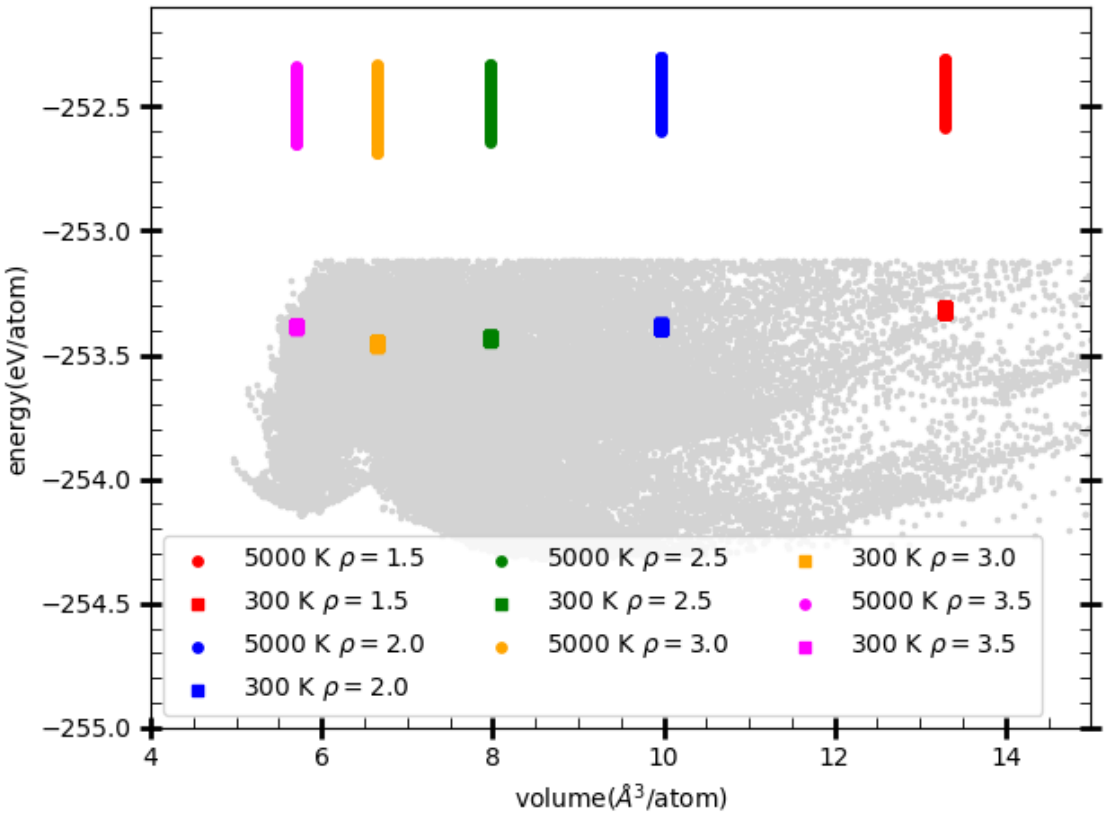

Figure B.3: The per atom energy-volume distribution of amorphous (300 K) and liquid (5000 K) structures sampled at different densities, compared to the training dataset (gray).

B.5 Distance Analysis

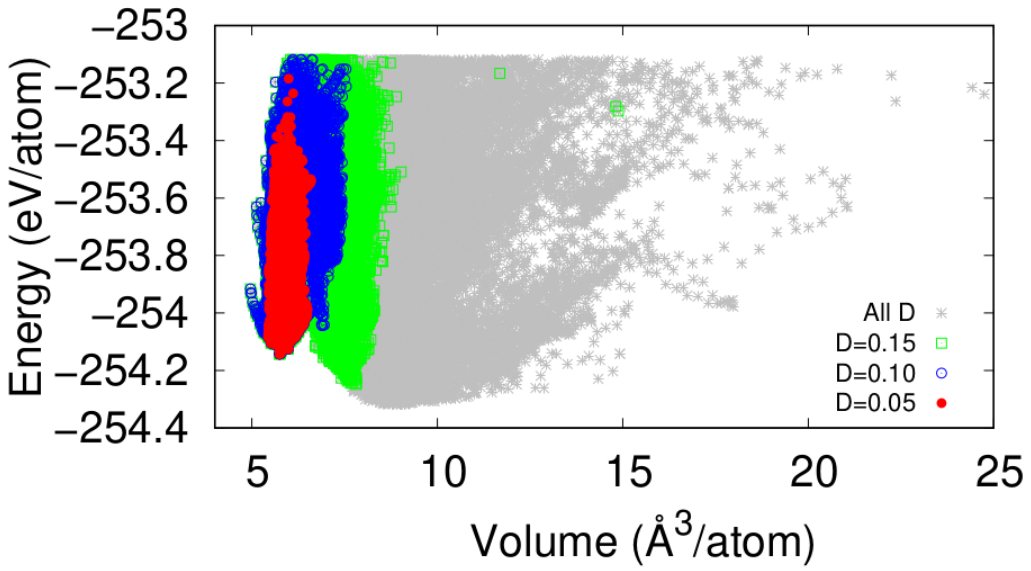

Figure B.4: Energy per atom as a function of volume for structures in the dataset, colored with respect to distance cutoff D from diamond.

Table B.2: RMSE for energy in meV/atom (top) and forces in eV/Å (bottom), for networks trained and validated on datasets of different distance to the reference phase, cubic diamond. T error is the average over the batch RMSE of the last 2500 training steps. The largest distance of any structure in the dataset to the reference phase is $D = 0.62$.

Train \ Validate	T error	All D	$D = 0.15$	$D = 0.10$	$D = 0.05$	$D_{12} = 0.05$
All D ($D = 0.62$)	22.0	**22.1**	20.9	15.2	7.7	7.3
$D = 0.15$	18.1	63.6	**17.8**	12.9	6.1	16.9
$D = 0.10$	8.7	162.3	52.2	**9.4**	4.3	76.3
$D = 0.05$	2.4	473.8	219.0	75.3	**2.5**	474.2
$D_{12} = 0.05$	2.6	174.3	88.3	52.0	2.7	**2.6**

Train \ Validate	T error	All D	$D = 0.15$	$D = 0.10$	$D = 0.05$	$D_{12} = 0.05$
All D ($D = 0.62$)	0.26	**0.27**	0.27	0.20	0.08	0.08
$D = 0.15$	0.24	0.49	**0.25**	0.18	0.08	0.15
$D = 0.10$	0.14	0.96	0.44	**0.15**	0.06	0.31
$D = 0.05$	0.05	1.67	1.05	0.55	**0.05**	0.31
$D_{12} = 0.05$	0.05	0.96	0.89	0.54	0.05	**0.05**

Bibliography

[1] J. R. Dahn, Tao Zheng, Yinghu Liu, and J. S. Xue. Mechanisms for lithium insertion in carbonaceous materials. *Science*, 270(5236):590–593, 1995.

[2] Elad Pollak, Baisong Geng, Ki-Joon Jeon, Ivan T. Lucas, Thomas J. Richardson, Feng Wang, and Robert Kostecki. The interaction of li+ with single-layer and few-layer graphene. *Nano Letters*, 10(9):3386–3388, 2010. PMID: 20677788.

[3] EunJoo Yoo, Jedeok Kim, Eiji Hosono, Hao-Shen Zhou, Tetsuichi Kudo, and Itaru Honma. Large reversible li storage of graphene nanosheet families for use in rechargeable lithium ion batteries. *Nano Letters*, 8(8):2277–2282, 2008.

[4] Guoxiu Wang, Xiaoping Shen, Jane Yao, and Jinsoo Park. Graphene nanosheets for enhanced lithium storage in lithium ion batteries. *Carbon*, 47(8):2049 – 2053, 2009.

[5] Jusef Hassoun, Francesco Bonaccorso, Marco Agostini, Marco Angelucci, Maria Grazia Betti, Roberto Cingolani, Mauro Gemmi, Carlo Mariani, Stefania Panero, Vittorio Pellegrini, and Bruno Scrosati. An advanced lithium-ion battery based on a graphene anode and a lithium iron phosphate cathode. *Nano Lett.*, 14(8):4901–4906, 2014.

[6] Bin Wang, Jaegeon Ryu, Sungho Choi, Gyujin Song, Dongki Hong, Chihyun Hwang, Xiong Chen, Bo Wang, Wei Li, Hyun-Kon Song, Soojin Park, and Rodney S. Ruoff. Folding graphene film yields high areal energy storage in lithium-ion batteries. *ACS Nano*, 12(2):1739–1746, 2018. PMID: 29350526.

[7] H. Sun, A. Varzi, V. Pellegrini, D.A. Dinh, R. Raccichini, A.E. Del Rio-Castillo, M. Prato, M. Colombo, R. Cingolani, B. Scrosati, S. Passerini, and F. Bonaccorso. How much does size really matter? exploring the limits of graphene as li ion battery anode material. *Solid State Communications*, 251:88 – 93, 2017.

[8] Eunseok Lee and Kristin A. Persson. Li absorption and intercalation in single layer graphene and few layer graphene by first principles. *Nano Lett.*, 12(9):4624–4628, 2012.

[9] Liu-Jiang Zhou, Z. F. Hou, and Li-Ming Wu. First-principles study of lithium adsorption and diffusion on graphene with point defects. *The Journal of Physical Chemistry C*, 116(41):21780–21787, 2012.

[10] Yasuharu Okamoto. Density functional theory calculations of lithium adsorption and insertion to defect-free and defective graphene. *The Journal of Physical Chemistry C*, 120(26):14009–14014, 2016.

[11] Chananate Uthaisar, Veronica Barone, and Juan E. Peralta. Lithium adsorption on zigzag graphene nanoribbons. *Journal of Applied Physics*, 106(11):113715, 2009.

[12] Ermias Girma Leggesse, Chi-Liang Chen, and Jyh-Chiang Jiang. Lithium diffusion in graphene and graphite: Effect of edge morphology. *Carbon*, 103:209–216, 2016.

[13] Xiaofeng Fan, W. T. Zheng, and Jer-Lai Kuo. Adsorption and diffusion of li on pristine and defective graphene. *ACS Applied Materials & Interfaces*, 4(5):2432–2438, 2012.

[14] Xiaofeng Fan, W. T. Zheng, Jer-Lai Kuo, and David J. Singh. Adsorption of single li and the formation of small li clusters on graphene for the anode of lithium-ion batteries. *ACS Applied Materials & Interfaces*, 5(16):7793–7797, 2013.

[15] Mingjie Liu, Alex Kutana, Yuanyue Liu, and Boris I. Yakobson. First-principles studies of li nucleation on graphene. *The Journal of Physical Chemistry Letters*, 5(7):1225–1229, 2014.

[16] Guangmin Yang, Xiaofeng Fan, Zhicong Liang, Qiang Xu, and Weitao Zheng. Density functional theory study of li binding to graphene. *RSC Adv.*, 6:26540–26545, 2016.

[17] Born M. and Oppenheimer J. R.. On the quantum theory of molecules. *Annelen der Physik*, 389(20):457–484, (1927).

[18] P. Hohenberg and W. Kohn. Inhomogeneous electron gas. *Phys. Rev.*, 136(3B):B864, (1964).

[19] W. Kohn and L. J. Sham. Self-consistent equations including exchange and correlation effects. *Phys. Rev.*, 140:A1133–A1138, Nov 1965.

[20] S. H. Vosko, L. Wilk, and M. Nusair. Accurate spin-dependent electron liquid correlation energies for local spin density calculations: a critical analysis. *Canadian Journal of Physics*, 58(8):1200–1211, 1980.

[21] J. P. Perdew and Alex Zunger. Self-interaction correction to density-functional approximations for many-electron systems. *Phys. Rev. B*, 23:5048–5079, May 1981.

[22] John P. Perdew and Yue Wang. Accurate and simple analytic representation of the electron-gas correlation energy. *Phys. Rev. B*, 45:13244–13249, Jun 1992.

[23] John P. Perdew, Kieron Burke, and Matthias Ernzerhof. Generalized gradient approximation made simple. *Phys. Rev. Lett.*, 77:3865–3868, Oct 1996.

[24] A. D. Becke. Density-functional exchange-energy approximation with correct asymptotic behavior. *Phys. Rev. A*, 38:3098–3100, Sep 1988.

[25] Chengteh Lee, Weitao Yang, and Robert G. Parr. Development of the colle-salvetti correlation-energy formula into a functional of the electron density. *Phys. Rev. B*, 37:785–789, Jan 1988.

[26] Riccardo Sabatini, Tommaso Gorni, and Stefano de Gironcoli. Nonlocal van der waals density functional made simple and efficient. *Phys. Rev. B*, 87:041108, Jan 2013.

[27] Walter Kohn, Yigal Meir, and Dmitrii E. Makarov. van der waals energies in density functional theory. *Phys. Rev. Lett.*, 80:4153–4156, May 1998.

[28] David C. Langreth and John P. Perdew. Exchange-correlation energy of a metallic surface: Wave-vector analysis. *Phys. Rev. B*, 15:2884–2901, Mar 1977.

[29] O. Gunnarsson and B. I. Lundqvist. Exchange and correlation in atoms, molecules, and solids by the spin-density-functional formalism. *Phys. Rev. B*, 13:4274–4298, May 1976.

[30] Stefan Grimme. Accurate description of van der waals complexes by density functional theory including empirical corrections. *Journal of Computational Chemistry*, 25(12):1463–1473, 2004.

[31] Stefan Grimme. Semiempirical gga-type density functional constructed with a long-range dispersion correction. *Journal of Computational Chemistry*, 27(15):1787–1799, 2006.

[32] Stefan Grimme, Jens Antony, Stephan Ehrlich, and Helge Krieg. A consistent and accurate ab initio parametrization of density functional dispersion correction (dft-d) for the 94 elements h-pu. *The Journal of Chemical Physics*, 132(15):154104, 2010.

[33] Alexandre Tkatchenko and Matthias Scheffler. Accurate molecular van der waals interactions from ground-state electron density and free-atom reference data. *Phys. Rev. Lett.*, 102:073005, Feb 2009.

[34] Axel D. Becke and Erin R. Johnson. Exchange-hole dipole moment and the dispersion interaction. *The Journal of Chemical Physics*, 122(15):154104, 2005.

[35] F. L. Hirshfeld. Bonded-atom fragments for describing molecular charge densities. *Theoretica chimica acta*, 44(2):129–138, Jun 1977.

[36] M. Dion, H. Rydberg, E. Schröder, D. C. Langreth, and B. I. Lundqvist. Van der waals density functional for general geometries. *Phys. Rev. Lett.*, 92:246401, Jun 2004.

[37] T. Thonhauser, Valentino R. Cooper, Shen Li, Aaron Puzder, Per Hyldgaard, and David C. Langreth. Van der waals density functional: Self-consistent potential and the nature of the van der waals bond. *Phys. Rev. B*, 76:125112, Sep 2007.

[38] Kyuho Lee, Éamonn D. Murray, Lingzhu Kong, Bengt I. Lundqvist, and David C. Langreth. Higher-accuracy van der waals density functional. *Phys. Rev. B*, 82:081101, Aug 2010.

[39] Valentino R. Cooper. Van der waals density functional: An appropriate exchange functional. *Phys. Rev. B*, 81:161104, Apr 2010.

[40] D. C. Langreth, M. Dion, H. Rydberg, E. Schröder, P. Hyldgaard, and B. I. Lundqvist. Van der waals density functional theory with applications. *International Journal of Quantum Chemistry*, 101(5):599–610, 2005.

[41] Andris Gulans, Martti J. Puska, and Risto M. Nieminen. Linear-scaling self-consistent implementation of the van der waals density functional. *Phys. Rev. B*, 79:201105, May 2009.

[42] Yingkai Zhang and Weitao Yang. Comment on "generalized gradient approximation made simple". *Phys. Rev. Lett.*, 80:890–890, Jan 1998.

[43] Oleg A. Vydrov and Troy Van Voorhis. Improving the accuracy of the nonlocal van der waals density functional with minimal empiricism. *The Journal of Chemical Physics*, 130(10):104105, 2009.

[44] Kyuho Lee, Éamonn D. Murray, Lingzhu Kong, Bengt I. Lundqvist, and David C. Langreth. Higher-accuracy van der waals density functional. *Phys. Rev. B*, 82:081101, Aug 2010.

[45] John P. Perdew and Wang Yue. Accurate and simple density functional for the electronic exchange energy: Generalized gradient approximation. *Phys. Rev. B*, 33:8800–8802, Jun 1986.

[46] Oleg A. Vydrov and Troy Van Voorhis. Nonlocal van der waals density functional: The simpler the better. *The Journal of Chemical Physics*, 133(24):244103, 2010.

[47] Oleg A. Vydrov and Troy Van Voorhis. Nonlocal van der waals density functional made simple. *Phys. Rev. Lett.*, 103:063004, Aug 2009.

[48] David C. Langreth and Bengt I. Lundqvist. Comment on "nonlocal van der waals density functional made simple". *Phys. Rev. Lett.*, 104:099303, Mar 2010.

[49] Oleg A. Vydrov and Troy Van Voorhis. Dispersion interactions from a local polarizability model. *Phys. Rev. A*, 81:062708, Jun 2010.

[50] Guillermo Román-Pérez and José M. Soler. Efficient implementation of a van der waals density functional: Application to double-wall carbon nanotubes. *Phys. Rev. Lett.*, 103:096102, Aug 2009.

[51] 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 Sclauzero, 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. *Journal of Physics: Condensed Matter*, 21(39):395502, 2009.

[52] 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 Jr, 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. *Journal of Physics: Condensed Matter*, 29(46):465901, 2017.

[53] Yusuf Shaidu, Emine Küçükbenli, and Stefano de Gironcoli. Lithium adsorption on graphene at finite temperature. *The Journal of Physical Chemistry C*, 122(36):20800–20808, 2018.

[54] Jang Wook Choi and Doron Aurbach. Promise and reality of post-lithium-ion batteries with high energy densities. *Nature Reviews Materials*, 1:16013, 2016.

[55] R. Yazami and Philip Touzain. A reversible graphite-lithium negative electrode for electrochemical generators. *Journal of Power Sources*, 9(3):365–371, 1983.

[56] Martin Winter, Jürgen O. Besenhard, Michael E. Spahr, and Petr Novak. Insertion electrode materials for rechargeable lithium batteries. *Advanced Materials*, 10(10):725–763, 1999.

[57] D. Guerard and A. Herold. Intercalation of lithium into graphite and other carbons. *Carbon*, 13(4):337–345, 1975.

[58] K. S. Novoselov, A. K. Geim, S. V. Morozov, D. Jiang, Y. Zhang, S. V. Dubonos, I. V. Grigorieva, and A. A. Firsov. Electric field effect in atomically thin carbon films. *Science*, 306(5696):666–669, 2004.

[59] P. E. Blöchl. Projector augmented-wave method. *Phys. Rev. B*, 50:17953–17979, Dec 1994.

[60] Andrea Dal Corso. Pseudopotentials periodic table: From h to pu. *Computational Materials Science*, 95:337–350, 2014.

[61] Hendrik J. Monkhorst and James D. Pack. Special points for brillouin-zone integrations. *Phys. Rev. B*, 13:5188–5192, Jun 1976.

[62] M. Methfessel and A. T. Paxton. High-precision sampling for brillouin-zone integration in metals. *Phys. Rev. B*, 40:3616–3621, Aug 1989.

[63] Kevin T. Chan, J. B. Neaton, and Marvin L. Cohen. First-principles study of metal adatom adsorption on graphene. *Phys. Rev. B*, 77:235430, Jun 2008.

[64] C. S. Praveen, Simone Piccinin, and Stefano Fabris. Adsorption of alkali adatoms on graphene supported by the au/ni(111) surface. *Phys. Rev. B*, 92:075403, Aug 2015.

[65] Y. Baskin and L. Meyer. Lattice constants of graphite at low temperatures. *Phys. Rev.*, 100:544–544, Oct 1955.

[66] A. M. Garay-Tapia, Aldo H. Romero, and Veronica Barone. Lithium adsorption on graphene: From isolated adatoms to metallic sheets. *Journal of Chemical Theory and Computation*, 8(3):1064–1071, 2012.

[67] Handan Yildirim, Alper Kinaci, Zhi-Jian Zhao, Maria K. Y. Chan, and Jeffrey P. Greeley. First-principles analysis of defect-mediated li adsorption on graphene. *ACS Applied Materials & Interfaces*, 6(23):21141–21150, 2014.

[68] Chananate Uthaisar and Veronica Barone. Edge effects on the characteristics of li diffusion in graphene. *Nano Letters*, 10(8):2838–2842, 2010.

[69] R. Car and M. Parrinello. Unified approach for molecular dynamics and density-functional theory. *Phys. Rev. Lett.*, 55:2471–2474, Nov 1985.

[70] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953.

[71] Shizukuni Yata, Hajime Kinoshita, Masatoshi Komori, Nobuo Ando, Takamitsu Kashiwamura, Tomoko Harada, Kazuyoshi Tanaka, and Tokio Yamabe. Structure and properties of deeply li-doped polyacenic semiconductor materials beyond c6li stage. *Synthetic Metals*, 62(2):153–158, 1994.

[72] David B. Laks, L. G. Ferreira, Sverre Froyen, and Alex Zunger. Efficient cluster expansion for substitutional systems. *Phys. Rev. B*, 46:12587–12605, Nov 1992.

[73] W. K. Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.

[74] Kristin Persson, Yoyo Hinuma, Ying Shirley Meng, Anton Van der Ven, and Gerbrand Ceder. Thermodynamic and kinetic properties of the li-graphite system from first-principles calculations. *Phys. Rev. B*, 82:125416, Sep 2010.

[75] Ruggero Lot, Franco Pellegrini, Yusuf Shaidu, and Emine Küçükbenli. Panna: Properties from artificial neural network architectures. *Computer Physics Communications*, 256:107402, 2020.

[76] Nongnuch Artrith and Jörg Behler. High-dimensional neural network potentials for metal surfaces: A prototype study for copper. *Phys. Rev. B*, 85:045439, Jan 2012.

[77] Patrick Rowe, Gábor Csányi, Dario Alfè, and Angelos Michaelides. Development of a machine learning potential for graphene. *Phys. Rev. B*, 97:054303, Feb 2018.

[78] Han Wang, Linfeng Zhang, Jiequn Han, and Weinan E. Deepmd-kit: A deep learning package for many-body potential energy representation and molecular dynamics. *Computer Physics Communications*, 228:178 – 184, 2018.

[79] Jörg Behler and Michele Parrinello. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.*, 98:146401, Apr 2007.

[80] Albert P. Bartók, Risi Kondor, and Gábor Csányi. On representing chemical environments. *Phys. Rev. B*, 87:184115, May 2013.

[81] M. Gastegger, L. Schwiedrzik, M. Bittermann, F. Berzsenyi, and P. Marquetand. wacsf—weighted atom-centered symmetry functions as descriptors in machine learning potentials. *The Journal of Chemical Physics*, 148(24):241709, 2018.

[82] Tian Xie and Jeffrey C. Grossman. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Phys. Rev. Lett.*, 120:145301, Apr 2018.

[83] Matthias Rupp, Alexandre Tkatchenko, Klaus-Robert Müller, and O. Anatole von Lilienfeld. Fast and accurate modeling of molecular atomization energies with machine learning. *Phys. Rev. Lett.*, 108:058301, Jan 2012.

[84] Katja Hansen, Franziska Biegler, Raghunathan Ramakrishnan, Wiktor Pronobis, O. Anatole von Lilienfeld, Klaus-Robert Müller, and Alexandre Tkatchenko. Machine learning predictions of molecular properties: Accurate many-body potentials and nonlocality in chemical space. *The Journal of Physical Chemistry Letters*, 6(12):2326–2331, 2015.

[85] Fleur Legrain, Jesús Carrete, Ambroise van Roekeghem, Stefano Curtarolo, and Natalio Mingo. How chemical composition alone can predict vibrational free energies and entropies of solids. *Chemistry of Materials*, 29(15):6220–6227, 2017.

[86] Masashi Tsubaki and Teruyasu Mizoguchi. Fast and accurate molecular property prediction: Learning atomic interactions and potentials with neural networks. *The Journal of Physical Chemistry Letters*, 9(19):5733–5741, 2018. PMID: 30081630.

[87] Jörg Behler. Atom-centered symmetry functions for constructing high-dimensional neural network potentials. *The Journal of Chemical Physics*, 134(7):074106, 2011.

[88] Linfeng Zhang, Jiequn Han, Han Wang, Roberto Car, and Weinan E. Deep potential molecular dynamics: A scalable model with the accuracy of quantum mechanics. *Phys. Rev. Lett.*, 120:143001, Apr 2018.

[89] Yufeng Huang, Jun Kang, William A. Goddard, and Lin-Wang Wang. Density functional theory based neural network force fields from energy decompositions. *Phys. Rev. B*, 99:064103, Feb 2019.

[90] Zhenwei Li, James R. Kermode, and Alessandro De Vita. Molecular dynamics with on-the-fly machine learning of quantum-mechanical forces. *Phys. Rev. Lett.*, 114:096405, Mar 2015.

[91] Volker L. Deringer and Gábor Csányi. Machine learning based interatomic potential for amorphous carbon. *Phys. Rev. B*, 95:094203, Mar 2017.

[92] Albert P. Bartók, Mike C. Payne, Risi Kondor, and Gábor Csányi. Gaussian approximation potentials: The accuracy of quantum mechanics, without the electrons. *Phys. Rev. Lett.*, 104:136403, Apr 2010.

[93] A.P. Thompson, L.P. Swiler, C.R. Trott, S.M. Foiles, and G.J. Tucker. Spectral neighbor analysis method for automated generation of quantum-accurate interatomic potentials. *Journal of Computational Physics*, 285:316 – 330, 2015.

[94] Mitchell A. Wood and Aidan P. Thompson. Extending the accuracy of the snap interatomic potential form. *The Journal of Chemical Physics*, 148(24):241721, 2018.

[95] Alexander V. Shapeev. Moment tensor potentials: A class of systematically improvable interatomic potentials. *Multiscale Modeling & Simulation*, 14(3):1153–1173, 2016.

[96] Yunxing Zuo, Chi Chen, Xiangguo Li, Zhi Deng, Yiming Chen, Jörg Behler, Gábor Csányi, Alexander V. Shapeev, Aidan P. Thompson, Mitchell A. Wood, and Shyue Ping Ong. Performance and cost assessment of machine learning interatomic potentials. *The Journal of Physical Chemistry A*, 124(4):731–745, 2020.

[97] J. S. Smith, O. Isayev, and A. E. Roitberg. ANI-1: an extensible neural network potential with dft accuracy at force field computational cost. *Chem. Sci.*, 8:3192–3203, 2017.

[98] Panna: Properties from artificial neural networks. `https://gitlab.com/PANNAdevs/panna`.

[99] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard,

Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.

[100] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv*, 1412.6980, 2014.

[101] S. Plimpton. Fast parallel algorithms for short-range molecular dynamics. *J. Comp. Phys.*, 117:1–19, 1995.

[102] Justin S. Smith, Olexandr Isayev, and Adrian E. Roitberg. ANI-1, a data set of 20 million calculated off-equilibrium conformations for organic molecules. *Scientific Data*, 4:170193 EP –, Dec 2017. Data Descriptor.

[103] Tobias Fink and Jean-Louis Reymond. Virtual exploration of the chemical universe up to 11 atoms of c, n, o, f: Assembly of 26.4 million structures (110.9 million stereoisomers) and analysis for new ring systems, stereochemistry, physicochemical properties, compound classes, and drug discovery. *Journal of Chemical Information and Modeling*, 47(2):342–353, 2007. PMID: 17260980.

[104] Rustam Z. Khaliullin, Hagai Eshet, Thomas D. Kühne, Jörg Behler, and Michele Parrinello. Graphite-diamond phase coexistence study employing a neural-network mapping of the ab initio potential energy surface. *Phys. Rev. B*, 81:100103, Mar 2010.

[105] Rustam Z. Khaliullin, Hagai Eshet, Thomas D. Kühne, Jörg Behler, and Michele Parrinello. Nucleation mechanism for the direct graphite-to-diamond phase transition. *Nature Materials*, 10:693 EP –, Jul 2011.

[106] J. Tersoff. Empirical interatomic potential for carbon, with applications to amorphous carbon. *Phys. Rev. Lett.*, 61:2879–2882, Dec 1988.

[107] Adri C. T. van Duin, Siddharth Dasgupta, Francois Lorant, and William A. Goddard. Reaxff: A reactive force field for hydrocarbons. *The Journal of Physical Chemistry A*, 105(41):9396–9409, 2001.

[108] Anthony K. Rappe and William A. Goddard. Charge equilibration for molecular dynamics simulations. *The Journal of Physical Chemistry*, 95(8):3358–3363, 1991.

[109] Emmanuel N. Koukaras, George Kalosakas, Costas Galiotis, and Konstantinos Papagelis. Phonon properties of graphene derived from molecular dynamics simulations. *Scientific Reports*, 5(1):12923, 2015.

[110] Mingjian Wen and Ellad B. Tadmor. Hybrid neural network potential for multilayer graphene. *Phys. Rev. B*, 100:195419, Nov 2019.

[111] Evgeny V. Podryabinkin and Alexander V. Shapeev. Active learning of linearly parametrized interatomic potentials. *Computational Materials Science*, 140:171 – 180, 2017.

[112] Nongnuch Artrith, Alexander Urban, and Gerbrand Ceder. Constructing first-principles phase diagrams of amorphous lixsi using machine-learning-assisted sampling with an evolutionary algorithm. *The Journal of Chemical Physics*, 148(24):241711, 2018.

[113] Linfeng Zhang, De-Ye Lin, Han Wang, Roberto Car, and Weinan E. Active learning of uniformly accurate interatomic potentials for materials simulation. *Phys. Rev. Materials*, 3:023804, Feb 2019.

[114] Patrick Rowe, Volker L. Deringer, Piero Gasparotto, Gábor Csányi, and Angelos Michaelides. An accurate and transferable machine learning potential for carbon. *The Journal of Chemical Physics*, 153(3):034702, 2020.

[115] Noam Bernstein, Gábor Csányi, and Volker L. Deringer. De novo exploration and self-guided learning of potential-energy surfaces. *npj Computational Materials*, 5(1):99, 2019.

[116] Stefano Curtarolo, Wahyu Setyawan, Shidong Wang, Junkai Xue, Kesong Yang, Richard H. Taylor, Lance J. Nelson, Gus L.W. Hart, Stefano Sanvito, Marco Buongiorno-Nardelli, Natalio Mingo, and Ohad Levy. Aflowlib.org: A distributed materials properties repository from high-throughput ab initio calculations. *Computational Materials Science*, 58:227 – 235, 2012.

[117] Juan J. de Pablo, Barbara Jones, Cora Lind Kovacs, Vidvuds Ozolins, and Arthur P. Ramirez. The materials genome initiative, the interplay of experiment, theory and computation. *Current Opinion in Solid State and Materials Science*, 18(2):99 – 117, 2014.

[118] Claudia Draxl and Matthias Scheffler. Nomad: The fair concept for big data-driven materials science. *MRS Bulletin*, 43(9):676–682, 2018.

[119] Colin W. Glass, Artem R. Oganov, and Nikolaus Hansen. Uspex—evolutionary crystal structure prediction. *Computer Physics Communications*, 175(11):713 – 720, 2006.

[120] Artem R. Oganov and Colin W. Glass. Crystal structure prediction using ab initio evolutionary techniques: Principles and applications. *The Journal of Chemical Physics*, 124(24):244704, 2006.

[121] Artem R. Oganov and Mario Valle. How to quantify energy landscapes of solids. *The Journal of Chemical Physics*, 130(10):104504, 2009.

[122] Mario Valle and Artem R. Oganov. Crystal fingerprint space – a novel paradigm for studying crystal-structure sets. *Acta Crystallographica Section A*, 66(5):507–517, Sep 2010.

[123] Steve Plimpton. Fast parallel algorithms for short-range molecular dynamics. *Journal of Computational Physics*, 117(1):1 – 19, 1995.

[124] Data repository for gaussian approximation potential. `http://www.libatoms.org/pub/Home/DataRepository`.

[125] thermo_pw: Ab-initio computation of material properties. `https://dalcorso.github.io/thermo_pw/`.

[126] Dario Alfè. Phon: A program to calculate phonons using the small displacement method. *Computer Physics Communications*, 180(12):2622 – 2633, 2009. 40 YEARS OF CPC: A celebratory issue focused on quality software for high performance, grid and novel computing architectures.

[127] Yanming Ma, Mikhail Eremets, Artem R. Oganov, Yu Xie, Ivan Trojan, Sergey Medvedev, Andriy O. Lyakhov, Mario Valle, and Vitali Prakapenka. Transparent dense sodium. *Nature*, 458(7235):182–185, 2009.

[128] Craig L. Bull, Giles Flowitt-Hill, Stefano de Gironcoli, Emine Küçükbenli, Simon Parsons, Cong Huy Pham, Helen Y. Playford, and Matthew G. Tucker. ζ-Glycine: insight into the mechanism of a polymorphic phase transition. *IUCrJ*, 4(5):569–574, Sep 2017.

[129] Nongnuch Artrith and Alexander Urban. An implementation of artificial neural-network potentials for atomistic materials simulations: Performance for tio2. *Computational Materials Science*, 114:135 – 150, 2016.

[130] Volker L. Deringer, Gábor Csányi, and Davide M. Proserpio. Extracting crystal chemistry from amorphous carbon structures. *ChemPhysChem*, 18(8):873–877, 2017.

[131] Muralikrishna Raju, P. Ganesh, Paul R. C. Kent, and Adri C. T. van Duin. Reactive force field study of li/c systems for electrical energy storage. *Journal of Chemical Theory and Computation*, 11(5):2156–2166, 2015. PMID: 26574418.

[132] John P. Perdew and Yue Wang. Accurate and simple analytic representation of the electron-gas correlation energy. *Phys. Rev. B*, 45:13244–13249, Jun 1992.

[133] H. J. McSkimin and P. Andreatch. Elastic moduli of diamond as a function of pressure and temperature. *Journal of Applied Physics*, 43(7):2944–2948, 1972.

[134] E. S. Zouboulis, M. Grimsditch, A. K. Ramdas, and S. Rodriguez. Temperature dependence of the elastic moduli of diamond: A brillouin-scattering study. *Phys. Rev. B*, 57:2889–2896, Feb 1998.

[135] E. J. Seldin and C. W. Nezbeda. Elastic constants and electron-microscope observations of neutron-irradiated compression-annealed pyrolytic and single-crystal graphite. *Journal of Applied Physics*, 41(8):3389–3400, 1970.

[136] Alexey Bosak, Michael Krisch, Marcel Mohr, Janina Maultzsch, and Christian Thomsen. Elasticity of single-crystalline graphite: Inelastic x-ray scattering study. *Phys. Rev. B*, 75:153408, Apr 2007.

[137] M. Mohr, J. Maultzsch, E. Dobardžić, S. Reich, I. Milošević, M. Damnjanović, A. Bosak, M. Krisch, and C. Thomsen. Phonon dispersion of graphite by inelastic x-ray scattering. *Phys. Rev. B*, 76:035439, Jul 2007.

[138] Daniel R. Cooper, Benjamin D'Anjou, Nageswara Ghattamaneni, Benjamin Harack, Michael Hilke, Alexandre Horth, Norberto Majlis, Mathieu Massicotte, Leron Vandsburger, Eric Whiteway, and Victor Yu. Experimental Review of Graphene. *ISRN Condensed Matter Physics*, 2012:1–56, 2012.

[139] Changgu Lee, Xiaoding Wei, Jeffrey W. Kysar, and James Hone. Measurement of the elastic properties and intrinsic strength of monolayer graphene. *Science*, 321(5887):385–388, 2008.

[140] Jae-Ung Lee, Duhee Yoon, and Hyeonsik Cheong. Estimation of young's modulus of graphene by raman spectroscopy. *Nano Letters*, 12(9):4444–4448, 2012. PMID: 22866776.

[141] P. Jacobson and S. Stoupin. Thermal expansion coefficient of diamond in a wide temperature range. *Diamond and Related Materials*, 97:107469, 2019.

[142] Monica Pozzo, Dario Alfè, Paolo Lacovig, Philip Hofmann, Silvano Lizzit, and Alessandro Baraldi. Thermal expansion of supported and freestanding graphene:

Lattice constant versus interatomic distance. *Phys. Rev. Lett.*, 106:135501, Mar 2011.

[143] P. J. Fallon, V. S. Veerasamy, C. A. Davis, J. Robertson, G. A. J. Amaratunga, W. I. Milne, and J. Koskinen. Properties of filtered-ion-beam-deposited diamondlike carbon as a function of ion energy. *Phys. Rev. B*, 48:4777–4782, Aug 1993.

[144] J. Schwan, S. Ulrich, H. Roth, H. Ehrhardt, S. R. P. Silva, J. Robertson, R. Samlenski, and R. Brenn. Tetrahedral amorphous carbon films prepared by magnetron sputtering and dc ion plating. *Journal of Applied Physics*, 79(3):1416–1422, 1996.

[145] B. Schultrich, H.-J. Scheibe, G. Grandremy, D. Drescher, and D. Schneider. Elastic modulus as a measure of diamond likeness and hardness of amorphous carbon films. *Diamond and Related Materials*, 5(9):914 – 918, 1996.

[146] B. Schultrich, H.-J. Scheibe, D. Drescher, and H. Ziegele. Deposition of superhard amorphous carbon films by pulsed vacuum arc deposition. *Surface and Coatings Technology*, 98(1):1097 – 1101, 1998. Papers presented at the Fifth International Conference on Plasma Surface Engineering.

[147] D. J. Evans and B. L. Holian. The nose–hoover thermostat. *The Journal of Chemical Physics*, 83(8):4069–4074, 1985.

[148] Lars Pastewka, Andreas Klemenz, Peter Gumbsch, and Michael Moseler. Screened empirical bond-order potentials for si-c. *Phys. Rev. B*, 87:205410, May 2013.

[149] Jeremy Bernstein, Arash Vahdat, Yisong Yue, and Ming-Yu Liu. On the distance between two neural networks and the stability of learning, 2020.

[150] M. A. Cusentino, M. A. Wood, and A. P. Thompson. Explicit multielement extension of the spectral neighbor analysis potential for chemically complex systems. *The Journal of Physical Chemistry A*, 124(26):5456–5464, 2020.

[151] Patrick Rowe, Volker L. Deringer, Piero Gasparotto, Gábor Csányi, and Angelos Michaelides. An accurate and transferable machine learning potential for carbon. *The Journal of Chemical Physics*, 153(3):034702, 2020.

[152] Liansheng Jiao, Zhenbang Liu, Zhonghui Sun, Tongshun Wu, Yuzhou Gao, Hongyan Li, Fenghua Li, and Li Niu. An advanced lithium ion battery based on a high quality graphitic graphene anode and a li[ni0.6co0.2mn0.2]o2 cathode. *Electrochimica Acta*, 259:48 – 55, 2018.

[153] So Fujikake, Volker L. Deringer, Tae Hoon Lee, Marcin Krynski, Stephen R. Elliott, and Gábor Csányi. Gaussian approximation potential modeling of lithium intercalation in carbon nanostructures. *The Journal of Chemical Physics*, 148(24):241714, 2018.

[154] Tsz Wai Ko, Jonas A. Finkler, Stefan Goedecker, and Jörg Behler. A fourth-generation high-dimensional neural network potential with accurate electrostatics including non-local charge transfer, arXiv.2009.06484, 2020.

[155] Tobias Morawietz, Vikas Sharma, and Jörg Behler. A neural network potential-energy surface for the water dimer based on environment-dependent atomic energies and charges. *The Journal of Chemical Physics*, 136(6):064103, 2012.

[156] S. Alireza Ghasemi, Albert Hofstetter, Santanu Saha, and Stefan Goedecker. Interatomic potentials for ionic systems with density functional accuracy based on charge densities obtained by a neural network. *Phys. Rev. B*, 92:045131, Jul 2015.

[157] Roohollah Hafizi, S. Alireza Ghasemi, S. Javad Hashemifar, and Hadi Akbarzadeh. A neural-network potential through charge equilibration for ws2: From clusters to sheets. *The Journal of Chemical Physics*, 147(23):234306, 2017.

[158] Somayeh Faraji, S. Alireza Ghasemi, Samare Rostami, Robabe Rasoulkhani, Bastian Schaefer, Stefan Goedecker, and Maximilian Amsler. High accuracy and transferability of a neural network potential through charge equilibration for calcium fluoride. *Phys. Rev. B*, 95:104105, Mar 2017.

[159] Aiichiro Nakano. Parallel multilevel preconditioned conjugate-gradient approach to variable-charge molecular dynamics. *Computer Physics Communications*, 104(1):59 – 69, 1997.

[160] Carlos Cárdenas, Paul Ayers, Frank De Proft, David J. Tozer, and Paul Geerlings. Should negative electron affinities be used for evaluating the chemical hardness? *Phys. Chem. Chem. Phys.*, 13:2285–2293, 2011.

[161] C. Ataca, E. Aktürk, S. Ciraci, and H. Ustunel. High-capacity hydrogen storage by metallized graphene. *Applied Physics Letters*, 93(4):043123, 2008.

[162] Fang Li and Jeffrey S. Lannin. Radial distribution function of amorphous carbon. *Phys. Rev. Lett.*, 65:1905–1908, Oct 1990.

[163] K. W. R. Gilkes, P. H. Gaskell, and J. Robertson. Comparison of neutron-scattering data for tetrahedral amorphous carbon with structural models. *Phys. Rev. B*, 51:12303–12312, May 1995.