On-the-Fly Training for Machine Learning Interaction Potentials

Carlos Rafael Salazar Letona
s21205751

February 1, 2023

Abstract

Machine learning (ML) methods have become more relevant in materials science over the last decade. They are now well recognized as effective methods for approximating extremely complex functions. Machine learning interaction potentials (MLIP) offer computation speeds close to empirical potentials, while having an accuracy close to the method used in training. In this project, an active learning approach was tested for the extension of gold databases with BCC and surface structures. The active learning approach used was Query by Committee (QbC), implemented in the lasp2 interface. The machine learning framework n2p2 was used for the creation of potentials and Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS) was used for exploring the phase space with NVT simulations. For selected structures, the vasp software was used to compute forces and energy using Density-Functional Theory (DFT), which where then added to the database. Potentials were then trained on-the-fly using lasp2 to measure the disagreement between the committee and then deciding when a certain structure should be added to the database. LAMMPS was then used to calculate surface energies for (100) and (111) structures, as well as BCC bulk modulus. These physical properties served as control parameters to test whether the potential was improving with on-the-fly training. It was found that the lasp2 method can improve a potential by doing on-the-fly training, however it is not always the case, as seen with BCC structures. For the (100) and (111) surfaces we were able to quickly recover the values predicted with a manually

created database. In the case of BCC, no clear improvement could be seen, and the error in bulk modulus seemed to behave arbitrarily, even when using the complete manually created database as a starting point.

Introduction

Material modelling at the atomic level is concerned with the calculation of the potential energy surface (PES). The PES is a multidimensional function that relates the energy of a molecule (or collection of molecules) and its geometry [1]. In materials science there are two main approaches to approximating the PES. On the one hand, electron-based methods are used for calculations at the electronic scales, they rely on quantum mechanics and are the most accurate models. However, they are also computationally expensive. On the other hand, *force field* methods are much faster, allowing for the simulation of larger systems, even reaching biological applications. But, they rely on empirical models, thus leading to a decrease in accuracy when compared to electron-based methods. In this project, a software for performing on-the-fly training of Machine Learning Interaction Potentials (MLIP) is tested.

Machine Learning Interaction Potentials

These are methods that aim to approach the accuracy of electronic-based methods, while having computation times similar to the atomic-based methods. This is done by using machine learning algorithms, such as neural network potentials [2], kernel based potentials [3], gaussian approximation potentials [4], among others.

High Dimensional Neural Network Potential

A High Dimensional Neural Network Potential (HDNNP) is a type of neural network potential (NNP) proposed by Behler and Parrinello [5]. It is created using Neural Networks (NN), consisting of a large amount of simple and highly interconnected processing units known as neurons [6]. Neurons are grouped in layers within the network, and each neuron is connected to all the neurons in the next layer. The connections between neurons are described by weights and biases. The HDNNP is constructed using *atomic* neural networks as building blocks, represented as the subnet S_i in fig 1.

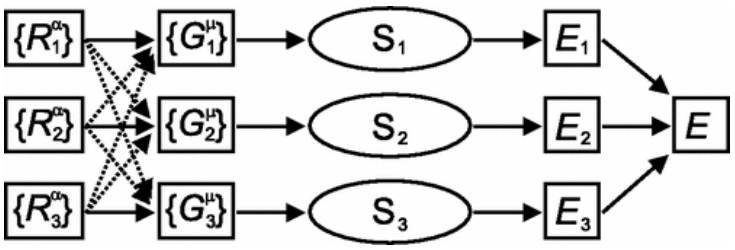

Figure 1: Diagram of a high dimensional neural network. R_i^α represents the cartesian coordinates of atom i, which are used to compute the symmetry function values G_i^μ used as input for the subnets S_i [5].

This complex network receives as input the cartesian coordinates of the atoms in the system. Symmetry functions are then computed, which are used to describe the system in a translational and rotational invariant manner [7]. The values of the symmetry functions are then used as input for the subnets S_i, giving as output the energy of the atom. The total energy of the system is then computed by a sum of the atomic energies, written mathematically as:

$$\begin{aligned} E_i &= f_a^2 \left[w_{01}^2 + \sum_{j=1}^{N_2} w_{j1}^2 f_a^1 \left(w_{0j}^1 + \sum_{\mu=1}^{N_1} w_{\mu j}^1 G_i^\mu \right) \right], \\ E &= \sum_{i=1}^{N_{atoms}} E_i \end{aligned} \tag{1}$$

where w_{ij}^k is the weight of the connection between neuron j in layer k and neuron i in layer $k-1$, w_{0j}^k is a bias, N_1 is the number of neurons in the input layer, N_2 is the number of neurons in the hidden layer, and f_a^k is the activation function for layer k. The atomic force can also be computed by calculating the derivative of this expression.

Query by Committee

Query by Committee (QbC) is an active learning strategy. In QbC a committee is formed, which consists of multiple machine learning models trained independently. This committee can then be used to compute an average of the predictions from its members, giving an improved accuracy compared to the individual predictions. Additionally, the disagreement of the members

in the committee provides access to an estimate of the generalization error. QbC consists in systematically improving the model by adding unlabeled data to the training set whose committee disagreement is high [8], [9].

Methodology

In this project a QbC method is tested for extending the training database of a HDNNP. The implementation of HDNNP used is n2p2 [10], which is used to approximate forces computed using Density Functional Theory (DFT) with the vasp software [11], [12], [13]. To perform molecular dynamics the LAMMPS software was used [14]. For the implementation of the on-the-fly training technique we developed the lasp2 (LAMMPS-vasp-n2p2) Python code [15].

Database

For our project, the purpose of using on-the-fly training is to automatically extend an existing database in order to train a better model. To be able to use our lasp2 method however, we need to have a starting database. A starting database can be created using a classical method that is known to be sufficient. With the auxiliary method we can explore the space of atomic structures. After the exploration, some structures are selected manually and the forces are computed using a more accurate method so that they may be added to the database. The database used in this project was generated by performing DFT calculations on structures obtained using an Embedded Atom Model (EAM) potential. This database consisted of 20 FCC, 25 BCC, 30 (100) surface, and 30 (111) surface structures of gold, shown in fig 2.

Simulation steps

The purpose of our lasp2 method is to automate the selection and calculation of structures for the database. We also eliminate the need for an auxiliary classical method to explore the phase space. This is achieved by training different models with the initial database, which will be used to create a committee. One of the models is then used to perform an NVT simulation using LAMMPS.

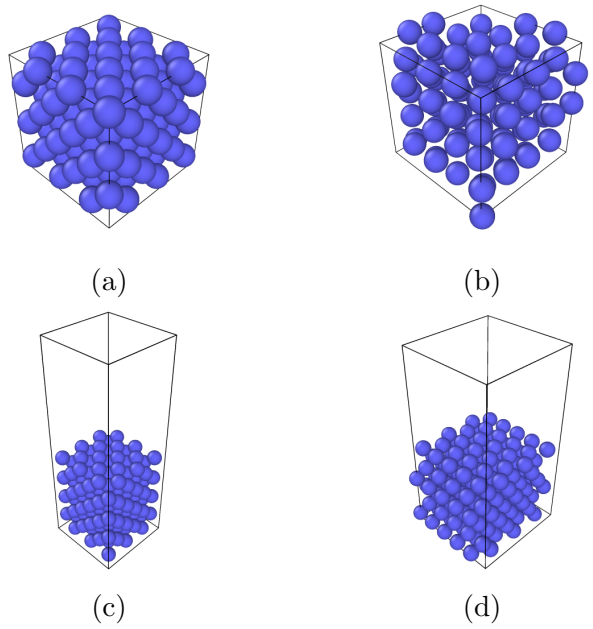

Figure 2: Figures (a) and (b) show the FCC and BCC bulk structures of gold, while figures (c) and (d) show the (100) and (111) surface structures of gold.

Since the initial database is expected to be limited, the disagreement of the committee is checked throughout the NVT simulation. If the disagreement is too high (greater than a user-defined threshold) for a given structure, it is likely that the model is not good for this type of structure. Therefore, the model would benefit from adding this structure to the database. Once the flag is activated (disagreement greater than threshold), the NVT simulation is stopped, and a DFT calculation is performed on the current structure. Then, the structure is added to the database and the models are retrained, finishing a lasp2 iteration. Afterwards, the NVT simulation is continued until the flag is activated again.

Disagreement

As discussed previously, the disagreement of the committee can be used as an indication for extrapolation (computing structures that are different than the training database). In our project we decided to use the standard deviation

as a measure of disagreement or dispersion of the predictions. This measure for the disagreement was chosen as it has been tested for similar applications [16]. In our case, it is implemented with a slight variation, as the maximum standard deviation is chosen as the value for disagreement. The disagreement γ is computed as:

$$\sigma_{F_\alpha} = \sqrt{\frac{\sum_{i=1}^{N}(F_{i_\alpha} - F_{\mu_\alpha})^2}{N}}, \qquad \gamma = \max(\sigma_{F_\alpha}) \tag{2}$$

where F_{i_α} is the force predicted by the i-th member of the committee for the atom α. F_{μ_α} is the average of the committee predictions for atom α. N is the number of members in the committee. Therefore, σ_{F_α} is the standard deviation of the committee predictions for atom α. Lastly, the disagreement measurement is equal to the maximum standard deviation.

An example of the dispersion throughout the NVT simulation can be seen in fig 3. Here, the disconnected lines represent different iterations of the lasp2 method. Once the dispersion goes over the threshold, training is performed and the NVT simulation restarts from the previous step.

Results

In order to test our method we selected subsets of the manually created database, and tried to recover their physical properties.

Extension of database to surface structures

For the first test, the database was limited to only 20 FCC structures. This subset was used to train the initial models.

Surface structure (100)

We performed the lasp2 method with an NVT simulation of a (100) surface structure at 400K. Each time the training flag is activated a lasp2 iteration occurs. At each lasp2 iteration, the database is increased by one (100) structure. For each iteration of the database, the surface energies for (100) and (111) were measured. The results are shown in fig 4 including the surface

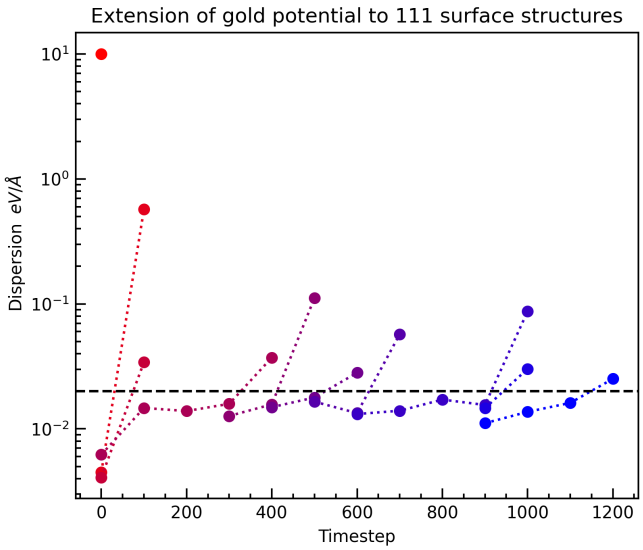

Figure 3: Dispersion for the different lasp2 iterations. The x axis represents the NVT simulation steps, and different plot lines represent different lasp2 iterations (red-oldest, blue-newest).

energy of the complete database as a straight dashed line. It can be seen that the error in surface energy compared to the DFT value is very high at iteration 0, when the database does not include any surface structure. Then, the error rapidly decreases for both (100) and (111) surfaces, implying that the (111) surface energy also benefits from adding (100) structures to the database. In the case of (100) surface energy, it manages to reach the original predictions of the complete database. However, this is not the case for (111) surface energy.

Surface structure (111)

The same process was then repeated using a (111) surface structure for the NVT simulation at 400K, and the results are shown in fig 5. This time, a (111) structure is added to the database for each lasp2 iteration. As in the previous case, the errors are high when no surface structure is included in the database. This time, however, the (111) surface energy reaches the original predictions from the complete database. This is just partially the case for

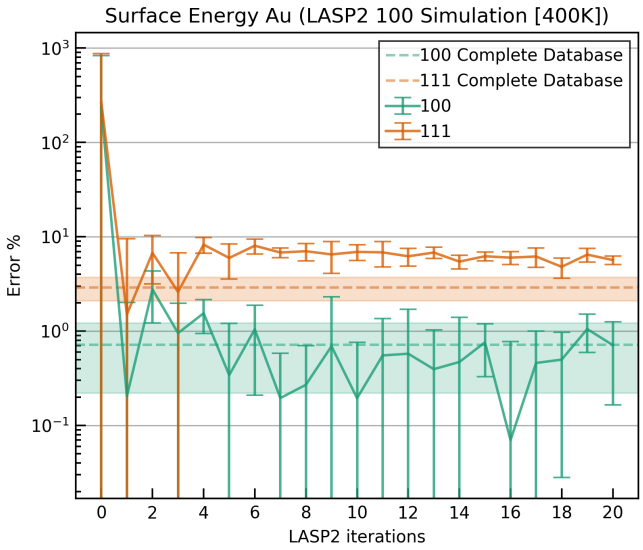

Figure 4: Error in surface energy for (100) and (111) surface structures with respect to the DFT value. The lasp2 simulation process was performed with an NVT simulation of a (100) surface structure. The value obtained using the original manually created database is shown as a dashed line with an error bar.

the (100) surface energy. The decrease in error is also fast, but slightly slower than in the (100) case.

Extension of database to BCC structures

Then, the same test was done for extending a limited database to BCC structures. This time, the bulk modulus of the BCC structure was used as a control variable. Different tests were performed with different subsets of the database, shown in fig 6. It is seen that for the BCC case there does not seem to be a decrease of the error with respect to DFT as a function of number of lasp2 iterations. Rather, the behaviour of the bulk modulus appears to be arbitrary, with few iterations approaching the original prediction from the complete database. However, it must be noted that the error of the complete database is still considerable, at $22.64\% \pm 1.42\%$. This might imply that the HDNNP in its current configuration is not sufficient for correctly

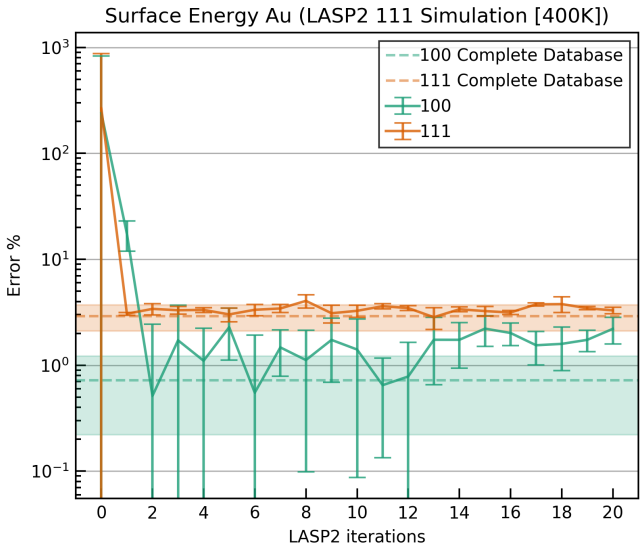

Figure 5: Error in surface energy for (100) and (111) surface structures with respect to the DFT value. The lasp2 simulation process was performed with an NVT simulation of a (111) surface structure. The value obtained using the original manually created database is shown as a dashed line with an error bar.

approximating the physical properties of BCC gold.

As an extra test, the lasp2 method was performed using the complete database for the initial model. The result is shown fig 7, where a similar behavior is observed. This further supports the hypothesis that the HDNNP struggles to compute physical properties of BCC gold.

Conclusion

This project serves as a proof of concept for the lasp2 method. However, the tests are not sufficient to assert its functionality. It is clearly shown that the calculation of physical properties of surface structures can be improved with this method. The error in surface energy shown in figs 4 and 5 quickly decreases, and reaches the predictions of the complete database just after a few iterations. In the case of BCC, however, the results are inconclusive, and

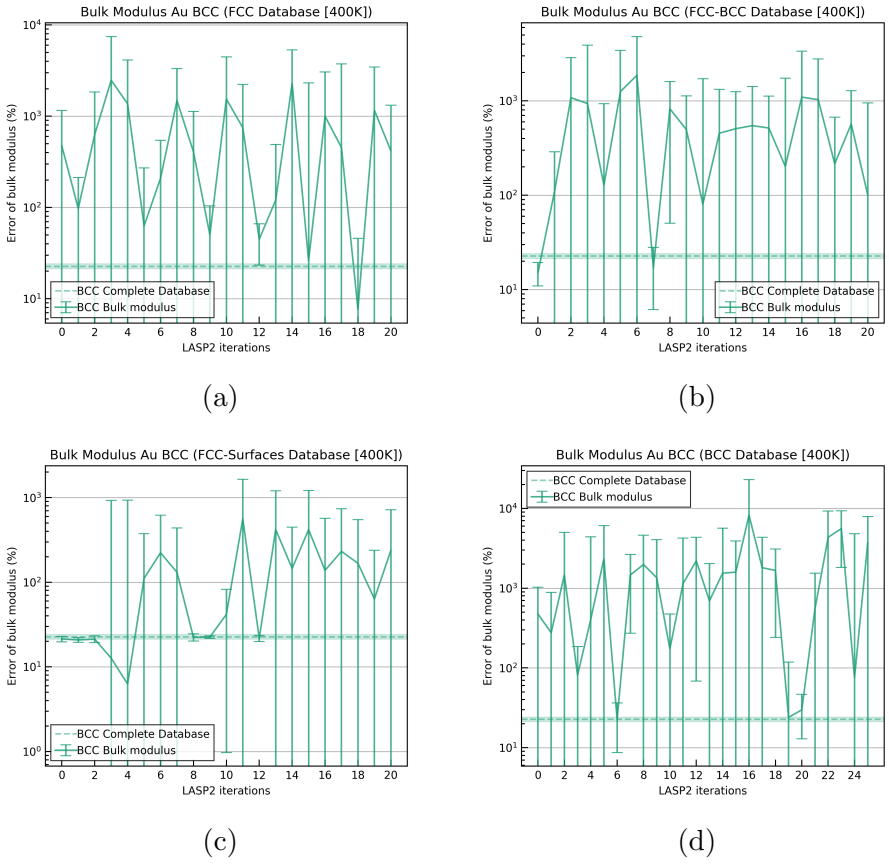

Figure 6: Error in bulk modulus for BCC structures with respect to the DFT value. The lasp2 simulation process was performed with an NVT simulation of a BCC structure. The value obtained using the original manually created database is shown as a dashed line with an error bar.

even when using larger subsets of the database, it was not possible to decrease the error in bulk modulus. This might be due to a limitation in the HDNNP model currently being used. However, in this early development stage, it is not possible to rule out technical problems with the code or the choice of parameters. Future work can be directed to testing the lasp2 method for different systems, and using different physical properties as control parameters. This could help rule out technical problems with the code.

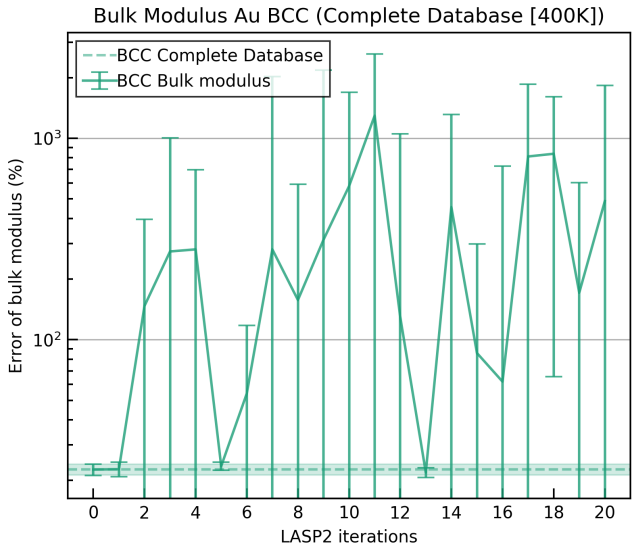

Figure 7: Error in bulk modulus for BCC structures with respect to the DFT value. The lasp2 simulation process was performed with an NVT simulation of a BCC structure. The value obtained using the original manually created database is shown as a dashed line with an error bar.

References

[1] Errol G. Lewars. "The Concept of the Potential Energy Surface". In: *Computational Chemistry: Introduction to the Theory and Applications of Molecular and Quantum Mechanics*. Cham: Springer International Publishing, 2016, pp. 9–49. ISBN: 978-3-319-30916-3. DOI: 10.1007/978-3-319-30916-3_2. URL: https://doi.org/10.1007/978-3-319-30916-3_2.

[2] Matti Hellström and Jörg Behler. "Neural Network Potentials in Materials Modeling". In: *Handbook of Materials Modeling: Methods: Theory and Modeling*. Ed. by Wanda Andreoni and Sidney Yip. Cham: Springer International Publishing, 2020, pp. 661–680. ISBN: 978-3-319-44677-6. URL: https://doi.org/10.1007/978-3-319-44677-6_56.

[3] Christoph Scherer, René Scheid, Denis Andrienko, and Tristan Bereau. "Kernel-Based Machine Learning for Efficient Simulations of Molecu-

lar Liquids". In: *Journal of Chemical Theory and Computation* 16.5 (2020), pp. 3194–3204. URL: https://doi.org/10.1021/acs.jctc.9b01256.

[4] Albert P. Bartók, Mike C. Payne, Risi Kondor, and Gábor Csányi. "Gaussian Approximation Potentials: The Accuracy of Quantum Mechanics, without the Electrons". In: *Phys. Rev. Lett.* 104 (13 Apr. 2010), p. 136403. URL: https://link.aps.org/doi/10.1103/PhysRevLett.104.136403.

[5] Jörg Behler and Michele Parrinello. "Generalized Neural-Network Representation of High-Dimensional Potential-Energy Surfaces". In: *Phys. Rev. Lett.* 98 (14 Apr. 2007), p. 146401. URL: https://link.aps.org/doi/10.1103/PhysRevLett.98.146401.

[6] Melody Y. Kiang. "Neural Networks". In: *Encyclopedia of Information Systems*. Ed. by Hossein Bidgoli. New York: Elsevier, 2003, pp. 303–315. ISBN: 978-0-12-227240-0. DOI: https://doi.org/10.1016/B0-12-227240-4/00121-0.

[7] Jörg Behler. "Atom-centered symmetry functions for constructing high-dimensional neural network potentials". In: *The Journal of Chemical Physics* 134.7 (2011), p. 074106. URL: https://doi.org/10.1063/1.3553717.

[8] H. S. Seung, M. Opper, and H. Sompolinsky. "Query by Committee". In: *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*. COLT '92. Pittsburgh, Pennsylvania, USA: Association for Computing Machinery, 1992, pp. 287–294. ISBN: 089791497X. DOI: 10.1145/130385.130417. URL: https://doi.org/10.1145/130385.130417.

[9] Anders Krogh and Jesper Vedelsby. "Neural Network Ensembles, Cross Validation, and Active Learning". In: *Advances in Neural Information Processing Systems*. Ed. by G. Tesauro, D. Touretzky, and T. Leen. Vol. 7. MIT Press, 1994. URL: https://proceedings.neurips.cc/paper/1994/file/b8c37e33defde51cf91e1e03e51657da-Paper.pdf.

[10] Andreas Singraber, mpbircher, Sam Reeve, David W.H. Swenson, Jérémy Lauret, and philippedavid. *CompPhysVienna/n2p2: Version 2.1.4*. Version v2.1.4. May 2021. URL: https://doi.org/10.5281/zenodo.4750573.

[11] G. Kresse and J. Hafner. "Ab initio molecular dynamics for liquid metals". In: *Phys. Rev. B* 47 (1 Jan. 1993), pp. 558–561. DOI: 10.1103/PhysRevB.47.558. URL: https://link.aps.org/doi/10.1103/PhysRevB.47.558.

[12] G. Kresse and J. Furthmüller. "Efficiency of ab-initio total energy calculations for metals and semiconductors using a plane-wave basis set". In: *Computational Materials Science* 6.1 (1996), pp. 15–50. ISSN: 0927-0256. DOI: https://doi.org/10.1016/0927-0256(96)00008-0. URL: https://www.sciencedirect.com/science/article/pii/0927025696000080.

[13] G. Kresse and J. Furthmüller. "Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set". In: *Phys. Rev. B* 54 (16 Oct. 1996), pp. 11169–11186. DOI: 10.1103/PhysRevB.54.11169. URL: https://link.aps.org/doi/10.1103/PhysRevB.54.11169.

[14] A. P. Thompson et al. "LAMMPS - a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales". In: *Comp. Phys. Comm.* 271 (2022), p. 108171. DOI: 10.1016/j.cpc.2021.108171.

[15] Carlos Rafael Salazar Letona. *LASP2Interface*. Version 1.0.0. Jan. 2023. DOI: 10.5281/zenodo.7592253. URL: https://lasp2interface.readthedocs.io/.

[16] Christoph Schran, Krystof Brezina, and Ondrej Marsalek. "Committee neural network potentials control generalization errors and enable active learning". In: *The Journal of Chemical Physics* 153.10 (2020), p. 104105. DOI: 10.1063/5.0016004. eprint: https://doi.org/10.1063/5.0016004. URL: https://doi.org/10.1063/5.0016004.