

KLIFF: A framework to develop physics-based and machine learning interatomic potentials ^{☆,☆☆}



Mingjian Wen¹, Yaser Afshar, Ryan S. Elliott, Ellad B. Tadmor^{*}

Department of Aerospace Engineering and Mechanics, University of Minnesota, Minneapolis, MN 55455, USA

ARTICLE INFO

Article history:

Received 11 August 2021

Received in revised form 1 October 2021

Accepted 30 October 2021

Available online 16 November 2021

Keywords:

Interatomic potentials

Machine learning

Uncertainty

OpenKIM

ABSTRACT

Interatomic potentials (IPs) are reduced-order models for calculating the potential energy of a system of atoms given their positions in space and species. IPs treat atoms as classical particles without explicitly modeling electrons and thus are computationally far less expensive than first-principles methods, enabling molecular simulations of significantly larger systems over longer times. Developing an IP is a complex iterative process involving multiple steps: assembling a training set, designing a functional form, optimizing the function parameters, testing model quality, and deployment to molecular simulation packages. This paper introduces the *KIM-based learning-integrated fitting framework* (KLIFF), a package that facilitates the entire IP development process. KLIFF supports both physics-based and machine learning IPs. It adopts a modular approach whereby various components in the fitting process, such as atomic environment descriptors, functional forms, loss functions, optimizers, quality analyzers, and so on, work seamlessly with each other. This provides a flexible framework for the rapid design of new IP forms. Trained IPs are compatible with the Knowledgebase of Interatomic Models (KIM) application programming interface (API) and can be readily used in major materials simulation packages compatible with KIM, including ASE, DL_POLY, GULP, LAMMPS, and QC. KLIFF is written in Python with computationally intensive components implemented in C++. It is parallelized over data and supports both shared-memory multicore desktop machines and high-performance distributed memory computing clusters. We demonstrate the use of KLIFF by fitting a physics-based Stillinger–Weber potential and a machine learning neural network potential for silicon. The KLIFF package, together with its documentation, is publicly available at: <https://github.com/openkim/kliff>.

Program summary

Program Title: KIM-based Learning-integrated Fitting Framework (KLIFF)

CPC Library link to program files: <https://doi.org/10.17632/fk77gs5b2d.1>

Licensing provisions: LGPL 2.1

Programming language: Python, C++

Supplementary material: Example scripts for the demonstrations, which are compatible with KLIFF v0.3.0.

Nature of problem: Development of a model called an *interatomic potential* (IP) representing the potential energy of a system of atoms based on their positions in space and species. This is a complex iterative process involving multiple steps: assembling a training set, designing a functional form, optimizing the function parameters, testing IP quality, and deployment of the fitted IP to molecular simulation packages.

Solution method: The fitting process is formulated as an optimization problem where a *loss function* characterizing the IP error over a training set is minimized to obtain the optimal fitting parameters. KLIFF is designed in a modular fashion providing the user with flexible access to different functional forms, loss functions, optimization algorithms, and analyzers for testing the quality of the fitted IP. KLIFF is built on

[☆] The review of this paper was arranged by Prof. Blum Volker.

^{☆☆} This paper and its associated computer program are available via the Computer Physics Communications homepage on ScienceDirect (<http://www.sciencedirect.com/science/journal/00104655>).

^{*} Corresponding author.

E-mail address: tadmor@umn.edu (E.B. Tadmor).

¹ Current address: Energy Technologies Area, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, United States.

the Knowledgebase of Interatomic Models (KIM) API standard, which enables immediate deployment of fitted IPs to major materials simulation packages that are compatible with KIM.

© 2021 Elsevier B.V. All rights reserved.

1. Introduction

Molecular simulations are a powerful computational technique for exploring material behavior and properties based on an understanding of the physics of bonding at the atomic scale [1]. This approach is used across the sciences with examples such as phase transition in crystals [2], protein folding [3], and thermal expansion and conductivity of layered 2D materials [4,5] to name just a few. At the core of any molecular simulation lies a description of the interactions between atoms that produces the forces governing atomic motion. First-principles approaches (e.g. density functional theory (DFT)) that involve solving the Schrödinger equation of quantum mechanics are most accurate, but due to hardware and algorithmic limitations, these approaches are limited to extremely small system sizes and time scales precluding the study of most systems of technological interest. For example, the supercell required to simulate a graphene bilayer with a 1.1° twist angle has more than 10,000 atoms, which is well beyond the capabilities of current first-principles approaches [5].

Interatomic potentials (IPs, also known as force fields) provide a classical alternative based on the Born–Oppenheimer approximation (BOA) [6]. Due to the large mass difference between nuclei and electrons, the BOA assumes that electrons instantaneously adapt to changes in nuclei positions adopting their ground state configuration – effectively decoupling nuclei and electron physics. This approximation is reasonable for many problems of interest in materials science and condensed-matter physics [1]. Consistent with the BOA, IPs treat atoms as classical particles without explicitly modeling the electrons, but strive to capture their influence on atomic nuclei in an effective manner. As such, IPs are computationally far less expensive than first-principles methods and can therefore be used to compute static and dynamic properties that are inaccessible to first-principles calculations [7–9]. In essence, an IP is a reduced-order model for the quantum-mechanical interaction of electrons and nuclei in a material through a parameterized functional form that depends only on the positions of the atomic nuclei (atoms hereafter).

Development of an IP is a complex iterative process involving multiple steps as shown in Fig. 1. (Refer back to this figure as you read the remainder of this section.) First, a dataset of experimental and/or first principles reference data must be assembled to which the IP will be fitted. When developing machine learning potentials, it is common practice to split the dataset into three parts: (1) a *training* set that is used to optimize the model parameters, (2) a *validation* set for fitting hyperparameters and monitor overfitting, and (3) a *test* set to assess the goodness of the fit.

Traditionally, the reference dataset contains material properties considered important for a given application, such as the cohesive energy, equilibrium lattice constant, and elastic moduli of given crystal phases to name a few. In recent years, many IPs adopt a *force-matching* scheme [10], in which the training set is augmented with the forces on atoms obtained by first-principles calculations for a large set of atomic configurations.² An advantage of this approach is that the issue of insufficient training data (particularly

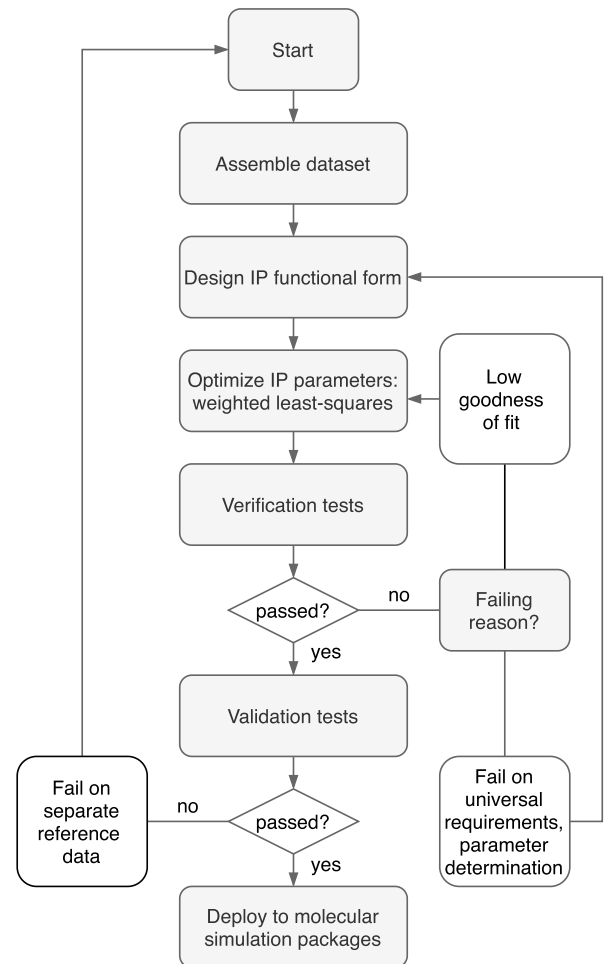


Fig. 1. Flowchart of the IP development process. Developing an IP involves four major steps: (1) assemble a set of reference data and design an IP functional form; (2) optimize IP parameters, typically carried out by minimizing a weighted least-squares loss function of the model predictions and the reference data; (3) assess the quality of the optimized model via verification and validation tests; and (4) deploy the model to molecular simulation packages. These steps can be iterative. When a model fails a verification test (e.g. by not satisfying a universal requirement, such as translational and rotational invariance, or by having a low goodness of fit on a test set) or fails a validation test (e.g. being unable to reproduce experimental material properties), it is necessary to return to earlier steps, make adjustments, and redo the fitting.

true for machine learning potentials) can be resolved because as many training data as needed can be readily generated.

Construction of a good reference dataset is critical for success. The fidelity of the IP for a given application hinges on including the appropriate physics in the dataset. It is also important to not swamp out rare configurations (such as transition states) that can have a disproportionate effect on material behavior. Dataset curation remains a difficult open problem and an area of active research [11].

Next an appropriate functional form has to be selected. Traditionally, the functional form of an IP was devised to represent the physics underlying the material system. One of the earliest examples is the pair potential developed by Lennard-Jones (LJ) in the 1920s to model van der Waals interactions in noble gases [12–14].

² These can be configurations associated with important structures or snapshots of the crystal as the atoms oscillate at finite temperature or through random perturbations.

The LJ potential includes an r^{-6} term (where r is the distance between atoms) that is based on a theoretical model for London dispersion, and an r^{-12} term meant to model repulsion due to Pauli exclusion. In the past century, a large number of physics-based potentials have been developed for a variety of ionic, metallic, and covalent systems [1]. A physics-based potential typically adopts a closed-form functional expression that is based on known physical or geometric aspects of bonding in the material. The functional forms of these IPs have become increasingly complex with an ever growing number of parameters.³

Devising the appropriate functional form to correctly capture the physics underlying the material system is arguably the most difficult task in developing a physics-based potential. It involves a mix of art and science as pointed out by Brenner [16]. This is largely alleviated by machine learning potentials [17–21], which have emerged in recent years and been shown to be highly effective for a spectrum of material systems ranging from organic molecules [19] to alloys [22]. Different from physics-based potentials, machine learning potentials are constructed by first transforming the atomic environment information in a large training set of first-principles results into vector representations (descriptors) and then training general-purpose regression algorithms on the atomic environment descriptors. In a machine learning potential, the regression algorithm contains no physics, but instead it attempts to “learn” the quantum mechanical Schrödinger equation directly from the training set of reference data. Properly tuned with a sufficiently dense training set, machine learning potentials have the advantage that, in principle, they can describe arbitrary bonding states and thus can achieve extremely high accuracy.

After the functional form has been selected (either physics-based or machine learning), the next step is to determine the values of the function parameters. This is typically formulated as a least-squares minimization problem by first constructing a *loss function* that quantifies the difference between the IP predictions and the reference values in the training set and then adjusting the parameters to reduce the loss function as much as possible. This can be challenging because IPs are nonlinear functions that are often “sloppy” in the sense that their predictions are insensitive to certain parameters or certain combinations of parameters [23,24]. These soft modes in parameter space can cause the minimization algorithms to fail to converge [9]. A solution is to use a minimization algorithm that moves along flat regions in parameter space more quickly (e.g. the geodesic Levenberg–Marquardt algorithm [25–27]), or better yet, to identify soft modes using a sensitivity analysis (e.g. a Fisher information based method [4]) and then apply a suitable model reduction.

Once an IP is trained, its quality must be assessed. This can be approached from a verification & validation (V&V) perspective. These terms are defined as [28]:

- **Verification:** The process of determining that a computational model accurately represents the underlying mathematical model and its solution.
- **Validation⁴:** The process of determining the degree to which a model is an accurate representation of the real world from the perspective of the intended uses of the model.

Verification for an IP includes satisfaction of universal requirements such as translational and rotational invariance (objectivity), permutation symmetry, forces returned by the IP correspond to the

³ For example, there are only two parameters in the LJ potential [14], whereas the ReaxFF [15] model developed for more complex systems has hundreds of adjustable parameters.

⁴ Note that the term *validation* is used differently in the V&V context than the validation set in machine learning mentioned above.

negative gradient of the energy, and so on. These are referred to as “verification checks” within the Knowledgebase of Interatomic Models (KIM) framework [29–32]. In addition, verification includes tests that assess the quality of the model in terms of the uncertainty in parameter determination, and the goodness of the fit using a test set as mentioned above.

The V&V notion of validation can be understood within the context of *transferability*, i.e. the ability of the IP to predict phenomena that it was not fit to reproduce. This includes prediction of material properties, computed by “KIM Tests” within the KIM framework [31], and predictions obtained through large-scale molecular simulations of real-world behavior. For example, the ability of IPs for carbon to reproduce the experimental structure of amorphous carbon [33].

As a general rule, physics-based potentials are better placed to exhibit transferability than machine learning potentials as long as the functional forms capture the requisite physics. For example an LJ potential fitted to the properties of an ideal gas provides a good approximation (within 10%) for the ground state crystal structure obtained by cooling the gas down to 0 K [1]. This is an impressive demonstration of transferability. In contrast, machine learning potentials have no physics beyond that in the training set (and possibly the descriptors). This means that a machine learning potential can only “transfer” to configurations that are close to what already exists in its training set.

Transferability can be included in the IP fitting process through a comparison of IP predictions with separate reference data. In cases where this fails, either the functional form needs to be extended for a physics-based potential, and/or the training set needs to be expanded for both physics-based and machine learning potentials. The training must then be redone.

Finally, once the IP fitting process is complete, the IP must be deployed to one or more molecular simulation packages of choice. Traditionally this is done on a code-by-code basis, which can be a time consuming and error prone process. If the IP class is already implemented in the code, then simply providing parameters may be enough – although even there things can go wrong. For example, for the REBO potential [34] implemented in LAMMPS [35], some of the parameters were not the ones presented in the original paper by Brenner et al. [34], but rather from the closely related AIREBO potential [36].⁵ In situations where an IP class is not available in a simulation code, the work involved in implementing it may be prohibitive. For example in the amorphous carbon study mentioned above [33] only IPs implemented in LAMMPS were tested, leaving out more than half of the possible IPs identified by the authors. The KIM application programming interface (API) [37] was designed to address this by creating a standard that allows a conforming IP to work seamlessly with any simulation code that supports it. The KIM API is supported by major materials simulation platforms including ASE [38,39], DL_POLY [40,41], GULP [42,43], LAMMPS [35,44,45], and QC [46,47].

This paper introduces the KIM-based learning-integrated fitting framework (KLIFF), a package that facilitates the entire IP development process described above. KLIF provides a unified Python interface to train both physics-based and machine learning potentials, and is constructed in modular fashion, making it easy to use and extend. It integrates closely with the KIM ecosystem for accessing IPs to train, testing trained IPs, and deploying trained IPs. The paper is structured as follows. Section 2 introduces two example IPs (one physics-based and the other machine learning) that will be trained later, and discusses the least-squares approach used to parameterize IPs. Section 3 presents KLIF features and capabilities. Implementation details of the code are outlined in Section 4.

⁵ This has been corrected in more recent implementations.

Section 5 presents a demonstration of using KLIF to fit the two IPs introduced in Section 2. The paper concludes in Section 6 with a summary.

2. Interatomic potentials

An IP is typically formulated as a parametric model that takes the positions of the atoms as its arguments and returns the potential energy,⁶

$$\mathcal{V} = \mathcal{V}(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N; \boldsymbol{\theta}), \quad (1)$$

where $\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N$ are the positions of a system of N atoms, and $\boldsymbol{\theta}$ denotes a set of fitting parameters associated with the IP functional form. An IP must be invariant with respect to rigid-body translation and rotation, inversion of space, and permutation of chemically equivalent species according to the laws of physics [1]. These symmetry requirements are typically intrinsic to the functional form of the IP. For example, if an IP is expressed in terms of distances between atoms, it automatically satisfies the requirements of translation, rotation and inversion invariance.

2.1. Physics-based potentials

The functional form of a physics-based potential is carefully devised to model the physics underlying the material system. For example, as discussed above, the LJ potential [12–14] provides a good model for van der Waals interactions in the noble gases, whereas for covalent systems more complex forms are required, such as bond-order potentials [15,34]. Here, we briefly review the three-body Stillinger–Weber (SW) potential for silicon [48] as an example.

The SW potential energy \mathcal{V} of a system consisting of N atoms has the form,

$$\mathcal{V} = \sum_{i=1}^N \sum_{j>i}^N \phi_2(r_{ij}) + \sum_{i=1}^N \sum_{j \neq i}^N \sum_{\substack{k>j \\ k \neq i}}^N \phi_3(r_{ij}, r_{ik}, \beta_{jik}), \quad (2)$$

where the two-body interaction takes the form

$$\phi_2(r_{ij}) = \epsilon \hat{A} \left[B \left(\frac{r_{ij}}{\sigma} \right)^{-p} - \left(\frac{r_{ij}}{\sigma} \right)^{-q} \right] \times \exp \left(\frac{1}{r_{ij}/\sigma - a} \right), \quad (3)$$

and the three-body term is

$$\phi_3(r_{ij}, r_{ik}, \beta_{jik}) = \epsilon \hat{\lambda} \left[\cos \beta_{jik} - \cos \beta^0 \right]^2 \times \exp \left(\frac{\hat{\gamma}}{r_{ij}/\sigma - a} + \frac{\hat{\gamma}}{r_{ik}/\sigma - a} \right), \quad (4)$$

in which $r_{ij} = \|\mathbf{r}_i - \mathbf{r}_j\|$ is the bond length between atoms i and j , β_{jik} is the bond angle formed by bonds $i-j$ and $i-k$ with the vertex at atom i . The parameters are $\epsilon, \hat{A}, B, p, q, \sigma, a, \hat{\lambda}, \hat{\gamma}$, and β^0 . The functional form is based on the lattice structure of bulk silicon shown in Fig. 2. The two-body term (Eq. (3)) models bond stretching and compression, and the three-body term (Eq. (4)) penalizes configurations away from the tetrahedral ground state structure of silicon.

The cutoff distance in the SW potential is implicitly defined as $r^{\text{cut}} = a\sigma$. This is not ideal from a potential fitting perspective. When fitting an IP, it is typical to fix the cutoff distance,

⁶ In general, IPs also depend on the species of the atoms. For notational simplicity, we limit our discussion to systems of a single atomic species. However, KLIF supports systems with arbitrary species.




Fig. 2. Bulk silicon crystallizes in a diamond cubic crystal structure in which each atom has four nearest neighbors forming the sp^3 hybridized tetrahedral structure.

and then adjust other parameters to minimize a loss function (discussed later in Section 2.3). For the standard form of SW, both a and σ must be fixed to set the cutoff. However, this adds an unnecessary constraint since two parameters are fixed instead of just the cutoff. If instead only a or σ are fixed (or neither), then the cutoff will vary during the fitting process. This can lead to failure of the optimization due to discontinuity in the loss function when neighbors enter or leave the cutoff sphere of an atom. In addition to the cutoff problem, another issue with the SW form is that ϵ is a redundant parameter that only scales the energy.

To avoid these pitfalls, Eqs. (3) and (4) are recast in a form in which all parameters are independent and the dependence on the cutoff radius is made explicit [4]. Let $A := \epsilon \hat{A}$, $\lambda := \epsilon \hat{\lambda}$, $\gamma := \sigma \hat{\gamma}$, and $r^{\text{cut}} := a\sigma$, we have

$$\phi_2(r_{ij}) = A \left[B \left(\frac{r_{ij}}{\sigma} \right)^{-p} - \left(\frac{r_{ij}}{\sigma} \right)^{-q} \right] \times \exp \left(\frac{\sigma}{r_{ij} - r^{\text{cut}}} \right), \quad (5a)$$

$$\phi_3(r_{ij}, r_{ik}, \beta_{jik}) = \lambda \left[\cos \beta_{jik} - \cos \beta^0 \right]^2 \times \exp \left(\frac{\gamma}{r_{ij} - r^{\text{cut}}} + \frac{\gamma}{r_{ik} - r^{\text{cut}}} \right). \quad (5b)$$

The new parameters are $A, B, p, q, \sigma, \lambda, \gamma$ along with the cutoff radius r^{cut} and the equilibrium bond angles β^0 . The SW model implemented in KIM [49] takes the form of Eqs. (5a) and (5b) instead of Eqs. (3) and (4).

2.2. Machine learning potentials

In contrast to physics-based potentials whose functional forms aim to capture the physics underlying the material system, machine learning potentials employ general-purpose regression models that interpolate across a dense training set of first principles energies and forces. Similar to a physics-based potential, a machine learning model returns the energy of an atom based on a finite neighborhood of atoms in its vicinity. Directly using the positions of an atom and its neighbors as input to the machine learning potential is ill-advised since this would require the model to learn the physical invariances of the IP [1,50], significantly increasing the complexity of the model and required training data. Instead, the atomic environment in terms of positions is transformed to a suitable “descriptor” vector representation that identically satisfies all invariances. For example two atomic environments that differ only by a rigid-body rotation would yield the same descriptor vector. Various descriptors have been developed to represent atomic environments, including the Coulomb matrix [19], symmetry functions [17,51], bispectrum [18,20,50], many-body tensor [52], and others [53]. As an example, we briefly review the symmetry functions approach, which is one of the earliest and most intuitive

representations. For a more detailed discussion, see for example Ref. [54].

The symmetry functions [17,51] are comprised of a set of two-body *radial functions* and a set of three-body *angular functions*. Specifically, the environment of atom i is characterized by three types of radial functions:

$$G_i^1 = \sum_{j \neq i} f_c(r_{ij}), \quad (6)$$

$$G_i^2 = \sum_{j \neq i} e^{-\alpha(r_{ij}-R_s)^2} f_c(r_{ij}), \quad (7)$$

$$G_i^3 = \sum_{j \neq i} \cos(\kappa r_{ij}) f_c(r_{ij}), \quad (8)$$

and two types of angular functions:

$$G_i^4 = 2^{1-\zeta} \sum_{j \neq i} \sum_{\substack{k > j \\ k \neq i}} (1 + \lambda \cos \beta_{jik})^\zeta e^{-\eta(r_{ij}^2 + r_{ik}^2 + r_{jk}^2)} \times f_c(r_{ij}) f_c(r_{ik}) f_c(r_{jk}), \quad (9)$$

$$G_i^5 = 2^{1-\zeta} \sum_{j \neq i} \sum_{\substack{k > j \\ k \neq i}} (1 + \lambda \cos \beta_{jik})^\zeta e^{-\eta(r_{ij}^2 + r_{ik}^2)} f_c(r_{ij}) f_c(r_{ik}), \quad (10)$$

where r_{ij} and β_{jik} are distance and angle as defined in Section 2.1, and α , R_s , κ , ζ , λ , and η are hyperparameters. The cutoff function f_c is given by

$$f_c(r) = \begin{cases} \frac{1}{2} \left[\cos\left(\frac{\pi r}{r^{\text{cut}}}\right) + 1 \right] & \text{for } r \leq r^{\text{cut}} \\ 0 & \text{for } r > r^{\text{cut}} \end{cases}, \quad (11)$$

where r^{cut} is the cutoff distance beyond which atoms do not contribute to the local environment.

The symmetry functions depend on both distances and angles, however since angles can be expressed in terms of distances through the law of cosines, the symmetry functions depend entirely on distances and are therefore invariant with respect to translation, rotation, and inversion of space [1]. The symmetry functions also satisfy the permutation symmetry requirement, because they are constructed by summation over all bond lengths and bond angles within the cutoff sphere and changing the summation order does not affect the results. One can select all the symmetry functions $G_i^1 \dots G_i^5$ to describe the atomic environment or a subset. As an example, we select one radial function and one angular function, G_i^2 and G_i^4 . The descriptor vector is comprised of distinct G_i^2 and G_i^4 values obtained for different choices of the hyperparameter sets $\{\alpha, R_s\}$ and $\{\lambda, \zeta, \eta\}$, respectively. The length of the descriptor vector is then equal to the total number of hyperparameter sets, $N_{G_i^2} + N_{G_i^4}$. (See the supplementary material for the hyperparameter sets for G_i^2 and G_i^4 used in Section 5.1.)

Many machine learning regression methods are suitable for constructing IPs including parametric linear regression and neural network (NN) models, nonparametric kernel ridge regression and Gaussian process models, and others [53]. Here, we discuss the NN model. In an NN potential, the total potential energy of a configuration consisting of N atoms is decomposed into the contributions of individual atoms

$$\mathcal{V} = \sum_{i=1}^N E_i, \quad (12)$$

where E_i is the energy of atom i , represented by an NN as shown in Fig. 3. The NN returns the energy E_i based on the positions of

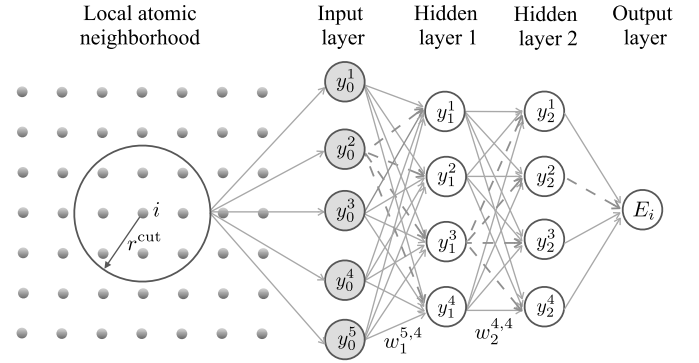


Fig. 3. Schematic representation of an NN potential to compute the atomic energy E_i . The NN consists of an input layer, two hidden layers and an output layer. The local atomic neighborhood information of atom i (all atoms within a sphere of radius r^{cut} around atom i) is transformed to descriptor vector with components y_0^j ($j = 1, 2, \dots$) that serves as the input to the NN. Each arrow connecting two nodes between adjacent NN layers represents a weight. The fully-connected NN becomes a dropout NN when some connections are cut (e.g. removing the dashed arrows). Biases and activation function are not shown in this plot. See text for explanation of the variables.

atom i and its neighbors up to a cutoff distance r^{cut} . The values y_0^1, y_0^2, \dots in the input layer are the components of the descriptor. Between the input layer and the energy output layer are so-called “hidden” layers that add complexity to the NN. In a fully-connected NN, each node in a hidden layer is connected to all the nodes in the previous layer and in the following layer. The value of node n in layer m is⁷

$$y_m^n = h \left(\sum_{n'} y_{m-1}^{n'} w_m^{n',n} + b_m^n \right), \quad (13)$$

where $w_m^{n',n}$ is the weight that connects node n' in layer $m-1$ and node n in layer m , b_m^n is the bias applied to node n of layer m , and h is an activation function (e.g. hyperbolic tangent) that introduces nonlinearity into the NN. In a more compact way, Eq. (13) can be written as $\mathbf{y}_m = h(\mathbf{y}_{m-1} \mathbf{W}_m + \mathbf{b}_m)$, where \mathbf{y}_m is a row vector of the node values in layer m , \mathbf{W}_m is a weight matrix, and \mathbf{b}_m is a row vector of the biases. For example, \mathbf{y}_1 and \mathbf{b}_1 are row vectors each with 4 elements and \mathbf{W}_1 is a 5×4 matrix for the NN shown in Fig. 3. Consequently, the atomic energy E_i represented in Fig. 3 can be expressed as

$$E_i = h[h[\mathbf{y}_0 \mathbf{W}_1 + \mathbf{b}_1] \mathbf{W}_2 + \mathbf{b}_2] \mathbf{W}_3 + \mathbf{b}_3. \quad (14)$$

The weights and biases are the fitting parameters in an NN potential: $\theta = \{\mathbf{W}_1, \mathbf{W}_2, \dots, \mathbf{W}_L, \mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_L\}$, where L is the number of layers (hidden and output).

2.3. Parameterization

Once an IP functional form is selected, the parameters must be determined. This is typically framed as a least-squares minimization problem where the IP parameters are adjusted to best match a training set of reference data obtained from experiments and/or first-principles calculations. For a training set of M configurations, the difference between the predictions of the IP and the reference data is quantified by a *loss function* defined as

⁷ The input layer and the output layer are indexed as the zeroth layer and third layer, respectively.

$$\mathcal{L}(\theta) = \frac{1}{2} \sum_{m=1}^M w_m^e \left[E(\mathbf{R}_m; \theta) - \hat{E}_m \right]^2 + \frac{1}{2} \sum_{m=1}^M w_m^f \|\mathbf{f}(\mathbf{R}_m; \theta) - \hat{\mathbf{f}}_m\|^2, \quad (15)$$

where $E(\mathbf{R}_m; \theta) \in \mathbb{R}$ and $\mathbf{f}(\mathbf{R}_m; \theta) = -(\partial E / \partial \mathbf{R})|_{\mathbf{R}_m} \in \mathbb{R}^{3N_m}$ are the energy and forces in configuration m obtained from an IP, \hat{E}_m and $\hat{\mathbf{f}}_m$ are the corresponding reference energy and forces for configuration m in the training set, with $\mathbf{R}_m \in \mathbb{R}^{3N_m}$ the concatenated coordinates of all atoms in configuration m and N_m the number of atoms in configuration m . The weights w_m^e and w_m^f are typically chosen to be inversely proportional to $(N_m)^2$, so that each configuration has an equal contribution to the loss function $\mathcal{L}(\theta)$. This prevents configurations with more atoms from dominating the optimization. For energy in units of eV and forces in units of eV/Å, these weights have units of eV⁻² and (eV/Å)⁻², respectively. Here, we only use energy and forces to construct the loss function, but in principle one can fit any physical property, such as the equilibrium lattice constants and elastic moduli of a ground state crystal structure. The objective then is to minimize the loss function in Eq. (15) with respect to θ to obtain the optimal set of IP parameters.

Simply minimizing Eq. (15) can lead to overfitting and thus low transferability of an IP. This is especially true for machine learning IPs due to the lack of physics in their functional forms and the large parameter space. Various techniques have been proposed to overcome this problem. One approach is to add regularization terms to the loss function to prevent overly complex results, for example an L_2 term of the form $\lambda \|\theta\|^2$ can be added, where λ is a hyperparameter that determines the regularization weight. Another widely used approach is early stopping [55], where model performance is monitored on a validation set and fitting is terminated when accuracy begins to degrade. There are also regularization techniques that are specific to certain types of models. For example, the dropout method [56,57] can be applied to NN potentials (see Section 3.2 for more on dropout).

3. Features and capabilities of KLIFF

A variety of software packages have been developed to develop IPs, including potfit [58,59], ænet [60], Amp [61], aPIP [62], atomcex [63], DeepMD-kit [64], GAP [18,65], MAISE [66], MLIP [67], PACE [68], PANNA [69], PyXtal_FF [70], RuNNer [17,71], SIMPLE-NN [72], among others. KLIFF shares many features with these packages, but is also distinguished by some capabilities described in this section that address the problems discussed in Section 1.

3.1. Integration with KIM

As indicated by the name, KLIFF is deeply integrated with the KIM ecosystem. (We note that the Potfit IP fitting framework is also compatible with KIM [9].)

First, KLIFF supports the training of IPs archived within the OpenKIM repository. An IP is called a *model* in KIM nomenclature, and a KIM *portable model* is an independent computer implementation of an IP that conforms to the KIM API portable model interface (PMI) standard.⁸ In practice portable models consist of a “model driver,” which implements an IP class (e.g. the embedded

⁸ KIM also supports a second type of model called a *simulator model*. While a portable model will work seamlessly with any simulation package that supports the KIM API/PMI standard, a simulator model only specifies how to set up and run a model that is implemented as an integrated part of a specific simulation package. KLIFF supports the fitting of portable models.

atom method (EAM) form) and a parameter set for a specific set of species. All content in the OpenKIM repository is archived subject to strict versioning and provenance control with digital object identifiers (DOIs) assigned. This makes it possible to access the exact IP used in a publication at a later date to reproduce the calculations or to conduct further fitting. A large number of physics-based and machine learning IPs are implemented as portable models and archived in the OpenKIM repository. These models are subjected to an editorial review process by the KIM Editor on acceptance to ensure quality control. Users of KLIFF can employ these models directly without having to implement them with significant savings in time and potential errors.

Second, IPs trained with KLIFF can be easily tested via OpenKIM. KLIFF can automatically generate models that are compatible with the KIM API, thus allowing a trained IP to run against KIM verification checks (VCs) and KIM tests [31]. As noted in Section 1, KIM VCs are programs that explore the integrity of an IP implementation. They check for programming errors (e.g. memory leak [73]), failures to satisfy required behaviors (e.g. inversion [74] and permutation [75] symmetries), and determine general characteristics of the IP functional form (e.g. are the forces returned by the model consistent with those obtained through numerical differentiation of the energy [76]). As opposed to KIM VCs, KIM tests check the accuracy of an IP by computing a variety of physical properties of interest to researchers, such as the stacking fault energy [77], elastic moduli [78], and linear thermal expansion coefficient [79] to name a few. The information provided by KIM VCs and KIM tests can save researchers a great deal of time by identifying limitations of an IP that can lead to subtle problems in simulations (e.g. poor convergence during energy minimization due to incorrect or discontinuous forces), and assisting in the selection of IPs by considering its predictions for relevant physical properties.

Third, IPs trained with KLIFF can be deployed via KIM. Traditionally, most IP development papers only report the functional form of the IPs and the associated parameters, without mentioning or providing a computer implementation. Recently developed machine learning potentials typically do provide computer implementations, but these are often standalone codes that cannot be used in major molecular simulation packages. This creates a significant barrier for the universal usability of IPs. By providing portable implementations, KIM addresses this issue, as well as enabling reproducibility.⁹ As mentioned above, KLIFF can automatically create IP models that are compatible with the KIM API, which enables the IP to work seamlessly with any KIM-compliant simulation package including ASE [38,39], DL_POLY [40,41], GULP [42,43], LAMMPS [35,44,45], and QC [46,47]. The final production IP can also be contributed to the OpenKIM repository for deployment as source and binary packages for major Mac, Linux and Windows platforms.

3.2. Uncertainty analysis

Historically, molecular simulation with IPs has been primarily viewed as a tool limited to providing qualitative insight. A key reason is that such simulations include multiple sources of uncertainty that are difficult to quantify, with the result that predictions obtained from the simulation typically lack robust confidence intervals [80]. A major source of uncertainty originates from the IPs themselves, since these are empirical functional forms fitted to experimental results and/or first-principles calculations. To make molecular simulation with IPs more reliable, it is imperative to quantify the intrinsic uncertainty of the IP and propagate it to

⁹ In some cases the same parameter file can lead to different results when used with different implementations of an IP, either in newer versions of the same code, or in different simulation packages. For example, see Ref. [8] for a discussion of this effect for tabulated EAM potentials.




Fig. 4. Schematic representation of the cost contours in the vicinity of the optimal parameters θ^* of an IP with two parameters θ_1 and θ_2 . The aspect ratio of the contours is determined by the eigenvalues λ_1 and λ_2 of the FIM. The diagonal elements of the inverse FIM $(\mathbf{F}^{-1})_{11}$ and $(\mathbf{F}^{-1})_{22}$ provide a lower bound on the variance of the parameters θ_1 and θ_2 , respectively.

the simulation results. This is an area that has not received much attention in the past. To address this limitation, KLIF provides functionality that enables uncertainty analysis of IPs.

As mentioned in Section 1, IPs are often “sloppy” [4,23,24] in that their predictions are insensitive to certain combinations of the parameters. This behavior can be quantified using the notion of a Fisher information matrix (FIM). The FIM provides a measure for the information in the training set on the parameters, which leads to an estimate for the precision with which the parameters can be determined [4]. For example, for the case where the loss function contains only forces (i.e. $w_m^e = 0$ in Eq. (15)), the FIM can be written as [4,24]:

$$\mathbf{F}(\theta) \propto \sum_{m=1}^M \left(\frac{\partial \mathbf{f}_m}{\partial \theta} \right)^T \left(\frac{\partial \mathbf{f}_m}{\partial \theta} \right), \quad (16)$$

where $\mathbf{f}_m \in \mathbb{R}^{3N_m}$ are the forces on atoms of configuration m (N_m is the number of atoms in configuration m), M is the number of configurations in the training set, and the superscript T denotes matrix transpose. The diagonal elements of the inverse FIM provide lower bounds on the variance of the parameters, known as the Cramér–Rao bound [81],

$$\text{Var}[\theta_i] \geq (\mathbf{F}^{-1})_{ii}. \quad (17)$$

The larger a diagonal element of the inverse FIM, the larger the lower bound on the variance for the corresponding parameter, indicating that the parameter is less well determined. As an illustrative example, we plot in Fig. 4 a schematic representation of the contours of the cost function in Eq. (15) for an IP with two parameters. Here, the two diagonal components of the inverse FIM $(\mathbf{F}^{-1})_{11}$ and $(\mathbf{F}^{-1})_{22}$ are nearly of the same magnitude, indicating that the two parameters θ_1 and θ_2 are equally determined in the fitting. If this were not the case and a diagonal component of the inverse FIM was much larger (an order of magnitude or more) than the others, then the IP parameter associated with this component is poorly determined. To address this, one could attempt to modify the IP functional form as discussed in Section 1 and shown in Fig. 1. The FIM also provides an upper bound on the uncertainty in a physical quantity of interest (QOI) due to small variations in IP parameters. A detailed discussion of such an analysis for the

thickness of monolayer MoS₂ can be found in Ref. [4]. The FIM in Eq. (16) is implemented in KLIF as an Analyzer (discussed in Section 4) using numerical differentiation.

The FIM analysis is well suited for physics-based potentials, which have dozens of parameters and each parameter plays a vital role. However, machine learning potentials are typically over-parameterized and the influence of a single parameter on the model performance is not large. Instead of parameter uncertainty, it is more important and useful to analyze the prediction uncertainty of a QOI (e.g. elastic moduli). A simple yet powerful approach to obtaining the QOI uncertainty is to construct an ensemble of IPs instead of a single best fit model. This can be done by either training different IPs using different initial guesses for the parameters or using different subsets of the training data. At the prediction stage, each individual model in the ensemble of N_p models is applied to compute the QOI P . The average

$$\bar{P} = \frac{1}{N_p} \sum_{i=1}^{N_p} P_i \quad (18)$$

is then used as the predictive mean for the QOI, and the standard deviation

$$\text{Std}[P] = \sqrt{\frac{1}{N_p - 1} \sum_{i=1}^{N_p} (P_i - \bar{P})^2} \quad (19)$$

as the uncertainty. The ensemble approach can be applied to any type of model, either physics-based [82,83] or machine learning potentials [84–87]. Although straightforward to train, it is computationally expensive since multiple models have to be trained to form the ensemble. For NN potentials, there is an alternative that is computationally less costly and performs equally well to the ensemble approach [88]. By removing some connections between layers (e.g. removing the dashed arrows for the NN shown in Fig. 3), a fully-connected NN is changed into a dropout NN [56,57]. It has been shown that training an NN with dropout (i.e. dropping different connections at each training step) approximates a Bayesian NN [89,90]. Consequently, a dropout NN possesses all the properties of a probabilistic Bayesian model, from which uncertainty information can be extracted. For dropout NN potentials [88], only one model needs to be trained at the training stage. At the prediction stage, it is essentially an ensemble model and can be used in a similar fashion: conduct multiple stochastic forward passes through the dropout NN (each time drop different connections) to obtain multiple samples of the QOI and then compute the average and standard deviation. KLIF supports the training of both ensemble and dropout NN potentials. The associated KIM DUNN model driver [91] allows molecular simulation codes to work with individual members in the ensemble and perform uncertainty quantification.

3.3. A wide range of support

By conforming to the KIM API, KLIF supports a wide range of IPs available through OpenKIM. At the time of this writing, the OpenKIM repository contains 35 model drivers, including widely used physics-based potentials such as Stillinger–Weber (SW) [48, 49], Tersoff [92–95], EDIP [96–99], and EAM [100–102] potentials among others. For machine learning potentials, KLIF currently supports the symmetry functions [17,51] and bispectrum [18,50] atomic environment descriptors. Interfacing with other descriptor libraries, such as DDescribe [103], is being explored. For machine learning regression algorithms, KLIF has its own implementation of simple algorithms (e.g. linear regression) and takes advantage of PyTorch [104] to build and train NN potentials. The NN model in

KLIFF wraps PyTorch so that the user interface appears the same as other models in KLIFF, but still retains the flexibility of PyTorch to create customizable NN structures and train with state-of-the-art deep learning techniques available through this package.

KLIFF provides an interface to many widely-used minimization algorithms for model training. As discussed in Section 2.3, the IP parameters are obtained by minimizing a loss function that quantifies the difference between IP predictions and the training set. The optimizer directly determines the values of the parameters and thus the quality of the IP. It is impossible to make a general statement about which optimizer is best, since this is problem-dependent, but some optimizers (e.g. the L-BFGS-B algorithm [105]) tend to work well for a wide range of problems. KLIFF supports the optimization algorithms in SciPy [106] and PyTorch [104]. The minimize module of `scipy.optimize` provides a large number of general-purpose minimization algorithms, and the `least_squares` module of `scipy.optimize` provides algorithms specific for nonlinear least-squares minimization problems. The optimizers in PyTorch are targeted for training NN models, including the stochastic gradient descent (SGD) method [107,108] and its variants such as the Adam method [109]. In addition, KLIFF also supports the geodesic Levenberg–Marquardt (LM) algorithm [25–27], which has been shown to work well for “sloppy” IPs whose predictions are insensitive to certain parameters or certain combinations of parameters.

3.4. Uniformity, modularity, and extensibility

KLIFF is designed to be as uniform, modular, and extensible as possible. It is implemented using an object-oriented programming (OOP) paradigm and provides a pure Python user interface. All the atomic environment descriptors, models, calculators, analyzers, etc. are subclassed from individual superclasses. A subclass only provides or modifies specific implementations of superclass methods when necessary, guaranteeing a uniform interface across subclasses. As mentioned in Section 3.3, KLIFF takes advantage of the optimization algorithms in SciPy [106] and PyTorch [104], as well as the geodesic LM algorithm to train models when minimization of a loss function is needed. Although vanilla SciPy, PyTorch, and LM have different APIs to call the optimization algorithms, KLIFF provides a unified interface that wraps them under the hood.

Extending KLIFF is straightforward. New descriptors, models, calculators, loss functions, optimization algorithms, analyzers, etc. can be seamlessly added to existing modules in KLIFF. For example, a new physics-based potential can be easily implemented by subclassing the KLIFF “Model” class, specifying the IP parameters, and then using Python to code the functional form of the IP. As a concrete example, we provide a Python code demonstrating how to implement the Lennard-Jones potential in the supplementary material. Other parts such as periodic boundary conditions handling are dealt with by KLIFF. The newly created model can then be used for training with any loss function and optimization algorithms that are available in KLIFF. To gain the benefits of KIM models discussed in Section 3.1, it is currently necessary to implement the IP as a separate code conforming to the KIM API. Future plans include the development of a general KIM model driver that will directly work with KLIFF IPs stored in a portable format.

3.5. Data parallelization

Computationally intensive KLIFF components, such as neighbor list generation and descriptor calculation, are internally implemented in C++. Even with this, the computational requirements can become quite demanding as the size of the training set increases. Fortunately, evaluation of the loss function Eq. (15) can be easily divided into independent sub-problems allowing for easy

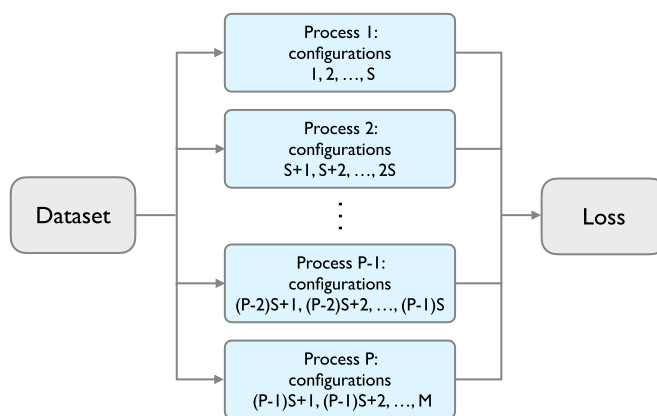


Fig. 5. Data parallelization scheme used by KLIFF. S is the number of configurations assigned to each process, and M is the total number of configurations.

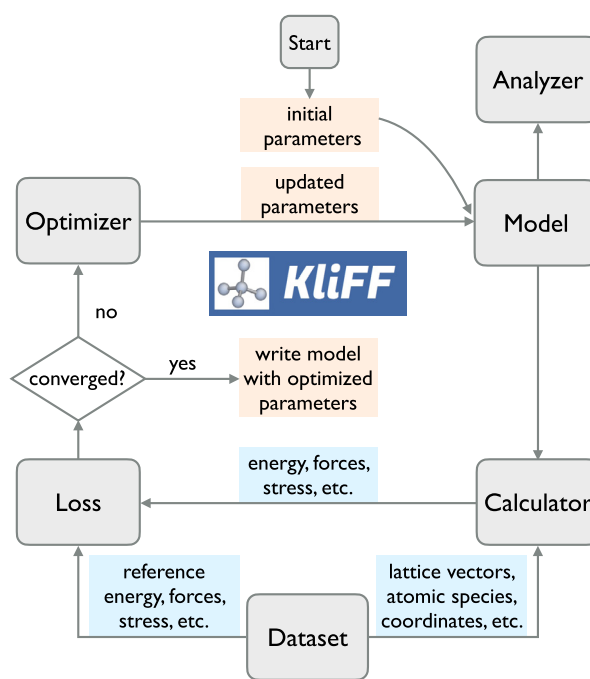


Fig. 6. Flowchart of the procedures of using KLIFF to train an IP.

parallelization. KLIFF adopts the parallelization over data scheme illustrated in Fig. 5. Atomic configurations in the dataset are distributed to different processes. Each process computes the sub-loss according to Eq. (15) for the configurations assigned to it, and the total loss is then obtained as the sum of the sub-losses from all the processes. KLIFF supports both OpenMP-style parallelism for shared-memory architectures, and MPI-style parallelism typical of high-performance computing clusters composed of multiple standalone machines connected by a network.

4. Implementation details: the KLIFF code

KLIFF is written primarily in Python with several computationally intensive components implemented in C++ accessible via Python bindings. As such, users interact with KLIFF through a pure Python interface. KLIFF is built in a modular fashion, as discussed in Section 3.4, with key modules `Dataset`, `Model`, `Calculator`, `Loss`, `Optimizer`, and `Analyzer`. A flowchart showing the interaction and information transfer between these modules for IP training is displayed in Fig. 6. The modules are described below.

4.1. Dataset module

A dataset is comprised of a set of atomic configurations, which provide the training data to optimize IP parameters or provide the test data to test the quality of an IP. An atomic configuration includes three vectors defining the simulation cell, flags to indicate whether periodic boundary conditions (PBCs) are applied along the cell vectors, the species and coordinates of all atoms in the configuration, and reference outputs. KLIF reads atomic configurations from extended XYZ files, with each configuration stored in a separate file. The reference outputs (energy, forces, and stress) associated with an atomic configuration are also read in from the extended XYZ file. The standard XYZ format only stores the number of atoms in a configuration and the species and coordinates of the atoms. The extended XYZ format allows for additional information to be stored, either in the second line via a series of `key=value` pairs (e.g. `PBC="T, T, T"` and `energy=1.2`) or in the body section by appending values (e.g. forces) to the coordinates. Internally, each atomic configuration and the reference outputs are associated with a `Configuration` object and a `Dataset` is essentially a set of `Configuration` objects.

4.2. Model module

The fitting process begins with the instantiation of a model (IP). Depending on the nature of the model, different operations can be applied. For KIM models and physics-based KLIF potentials, KLIF can provide information on what parameters are available for fitting, together with a description of each parameter and the data structure and data type of each parameter. Based on this information, a user can select the parameters to fit and specify initial values or use defaults. Lower and upper bounds on parameter values can also be provided to restrict it to a range. For an NN model, the descriptor representation of an atomic environment, which serves as the input to the NN model, must be defined. Then the NN can be constructed using an arbitrary number of layers, nodes per layer, and activation functions. Unlike physics-based models, KLIF automatically initializes the parameters in the network. For example, the He initializer [110] is used to initialize the weights and biases in Eq. (13). Other default choices are made by KLIF based on the authors' physical understanding and experience to make it easier for users to develop machine learning potentials without having to master subtle aspects of machine learning training. For example, in a standard dropout approach, different NN connections would be removed for each atom in a configuration (see Section 3.2). However, KLIF defaults to a native dropout scheme that removes the same NN connections for all atoms in a configuration. This ensures that atoms with identical environments (e.g. all atoms in an ideal silicon crystal) will have the same atomic energy, forces, and other properties. Users can overwrite default choices, for example, by selecting the native PyTorch dropout instead of KLIF's native implementation.

4.3. Calculator module

The created model is attached to a *calculator* that computes the predictions corresponding to the reference outputs for the atomic configurations in the training set. The native KLIF calculator supports the evaluation of energy, forces, and stress. If a property other than these is to be fitted, a new calculator needs to be implemented. A new calculator can wrap any KIM compliant molecular simulation package to compute the property with the given model in a similar fashion to ASE calculators [38,39].

4.4. Loss module

The predictions computed by the calculator and the corresponding reference output values stored in the training set are then used to construct a loss function (e.g. Eq. (15)) that quantifies the difference between the model predictions and the references. A weight can be assigned to each configuration, so that "important" configurations are emphasized more during optimization. If the available loss functions in KLIF do not satisfy a specific need, a user-defined loss function can be added.

4.5. Optimizer module

The optimization process involves minimization of the loss function with respect to the IP parameters until specified stopping criteria are satisfied, such as reducing the loss function value below a tolerance or reaching a maximum allowed number of minimization steps.

The optimizers supported by KLIF can be broadly categorized in two classes: *batch optimizers* and *mini-batch optimizers*. The former (e.g. the L-BFGS-B and geodesic LM methods) typically require the evaluation of the entire training set at each minimization step, whereas the latter (e.g. the SGD and Adam methods) only use a subset of the training set at a time. Batch optimizers guarantee a monotonic decrease of the loss throughout the minimization process and typically yield smaller final loss values compared with mini-batch optimizers. Mini-batch optimizers become advantageous for very large training sets (typical of machine learning potentials) where evaluation of the entire training set becomes prohibitive due to memory and/or computing constraints. For NN models that contain a large number of parameters, SGD-based optimizers can typically find a reasonable solution in parameter space that minimizes the loss to a certain level. By default, KLIF uses an L-BFGS-B optimizer for physics-based potentials, which typically have relatively small numbers of parameters and small training sets, and an SGD-based Adam optimizer for NN potentials, which have many parameters and very large training sets. The user can overwrite this default and select a preferred optimizer.

Once the optimization is completed, the fitted IP can be written out as a KIM model that conforms to the KIM API, which can then be run against KIM VCs and KIM tests or be used with any KIM-compliant simulations codes as discussed in Section 3.1. Generated KIM models, can be uploaded to <https://openkim.org> to receive a DOI and make the model available to the broader research community. Also, the model can be attached to an *Analyzer* to carry out post-processing analysis, such as computing the FIM discussed in Section 3.2 and computing the root mean square errors of energy and forces for a test set.

4.6. Command line tool

KLIF provides a command line tool called `kliff` that facilitates the execution of many common tasks. For example, query a physics-based potential for available parameters that can be optimized and their associated metadata, print a synopsis of the atomic configurations in the dataset, or split a dataset into multiple subsets. Once installed, executing `"kliff --help"` in the terminal will list the commands, their arguments, and help information.

5. Demonstration

KLIF has been extensively tested through the development of multiple IPs, including an SW potential for two-dimensional molybdenum disulfide [4], an interlayer potential for multilayer graphene [111], a hybrid NN potential for multilayer graphene [5],

Table 1

Summary of SW parameters obtained by minimizing the loss function and the preset parameters.

| Parameter | Value | Parameter | Value |
|-----------|----------------|------------------|----------------|
| A | 15.46588611 eV | B | 0.61032816 |
| p | 4 | q | 0 |
| σ | 2.05971554 Å | λ | 65.46736831 eV |
| γ | 2.71009995 Å | r^{cut} | 3.77118 Å |
| β^0 | 109.47° | | |

and a dropout uncertainty NN potential (DUNN) to quantify uncertainty in molecular simulations [88]. In this section we present examples demonstrating the use of KLIF in training an SW potential and an NN potential for silicon. The functional forms of the two IPs are described in Section 2.

5.1. Parameterization

The training set is comprised of the energies and forces for 2513 configurations of silicon in the diamond cubic crystal structure. This includes configurations with compressed and stretched cells and random perturbations of atoms, as well as configurations drawn from a molecular dynamics trajectory at a temperature of 300 K. Since this is only a demonstration, instead of using first-principles calculation or experimental data, the configurations were generated using the EDIP model [96–99]. The dataset is provided in the supplementary material.

The SW potential has seven parameters, A , B , p , q , σ , λ , γ , along with the cutoff radius r^{cut} and the equilibrium angle β^0 . The cutoff radius is set to $r^{\text{cut}} = 3.77118$ Å, as used by Stillinger and Weber [48], and the equilibrium angle is set to the tetrahedral angle of the ideal cubic diamond structure, $\beta^0 = 109.47^\circ$. Following most SW parameterizations [4,48,112], the parameters p and q are set to 4 and 0, respectively. The values of the remaining parameters are obtained by minimizing the loss function in Eq. (15) using the geodesic LM algorithm [25–27]. The energy and force weights are set to $w_m^e = 1/(N_m)^2$ and $w_m^f = 10/(N_m)^2$. A larger force weight is used to better reproduce the phonon dispersions discussed in Section 5.2. One exception is that the energy weight is set to $w_m^e = 10/(N_m)^2$ for configurations that have an ideal cubic diamond structure at different lattice parameters. The increased weight ensures that these configurations are not underrepresented in the fitting since their force terms in Eq. (15) are identically zero (regardless of the IP parameters) due to the symmetry of the underlying structure. The optimal parameter set identified by this process and the preset parameters are listed in Table 1.

For the NN potential, we employ the G_i^2 and G_i^5 symmetry functions (Eqs. (7) and (10)) as the descriptors for characterizing atomic environments. The hyperparameters α and R_s in Eq. (7) and ζ , λ , and η in Eq. (10) are provided in the supplementary material. The cutoff in Eq. (11) is set to $r^{\text{cut}} = 3.5$ Å to include only nearest-neighbor interactions. A challenging aspect of training an NN, which is also a source of the power and flexibility of the method, is that it is up to the developer to select the number of descriptor terms to retain, the number of hidden layers, the number of nodes within each hidden layer (which need not be the same), and the activation function. It is also possible to create different connectivity scenarios between layers. Here we have opted for simplicity and adopted a fully-connected network with the same number of nodes in each hidden layer. The number of hidden layers and the number of nodes in each hidden layers are determined through a grid search and are listed in Table 2. The activation function h is taken to be the commonly used hyperbolic tangent function, $\tanh(x) = (e^x - e^{-x})/(e^x + e^{-x})$.

The NN potential parameters are obtained by minimizing the loss function Eq. (15). The energy weight w_m^e and forces weight

Table 2

Summary of parameters in the NN potential and hyperparameters that define the NN structure.

| | |
|----------------------------------|----------------------------|
| number of hidden layers | 3 |
| number of nodes in hidden layers | 10 |
| cutoff r^{cut} | 3.5 Å |
| activation function h | tanh |
| descriptor hyperparameters | see supplementary material |

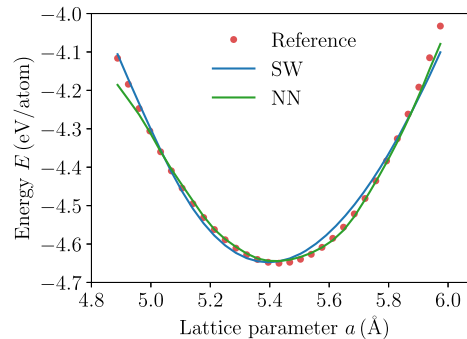


Fig. 7. Potential energy of diamond cubic silicon as a function of the lattice parameter predicted by the trained SW and NN potentials along with the EDIP reference data.

Table 3

Cohesive energy (absolute value of the minimum of the potential energy versus lattice parameter curve) and equilibrium lattice constant for diamond cubic silicon computed using the EDIP potential (taken as the reference) and the SW and NN potentials (with errors relative to EDIP given in parentheses).

| Potential | E_{coh} [eV/atom] | a_0 [Å] |
|-----------|----------------------------|--------------|
| EDIP | 4.650 | 5.43 |
| SW | 4.647 (0.06%) | 5.39 (0.74%) |
| NN | 4.645 (0.1%) | 5.42 (0.18%) |

w_m^f are the same as those used for the SW potential. The minimization is carried out using the Adam optimizer [109] with a learning rate of 0.001. As discussed in Section 4, to accelerate the training process a mini-batch technique [113] is employed with a batch size of 100 configurations at each minimization step for a total of 2000 epochs.¹⁰

The scripts used to train the SW and NN potentials are provided in the supplementary material.

5.2. Testing the trained potentials

To test the fitted SW and NN potentials, we applied them to study energetic and vibrational properties of silicon in the diamond cubic crystal structure. As discussed in Section 3.1, IPs trained by KLIF can be exported in a form compatible with the KIM API, which allows them to be used directly with a variety of major molecular simulation packages, such as LAMMPS [35,44,45]. The tests described in this section were carried out using LAMMPS.

First, we investigate the cohesive energy versus lattice parameter for ideal cubic diamond silicon (see structure in Fig. 2). The fitted SW and NN potentials are compared with the EDIP reference data in Fig. 7. Both potentials reproduce the equilibrium state well as seen in Table 3, however the NN potential with its flexible functional form is able to follow the reference data more closely

¹⁰ An epoch is one complete pass over the dataset. For example, if a dataset includes 50 configurations and a mini-batch size of 10 configurations is used, then one epoch consists of 5 minimization steps.




Fig. 8. Potential energy of diamond cubic silicon as a function of the lattice parameter predicted by the DUNN potential. (a) Predictive mean and uncertainty of the energy by DUNN, where the uncertainty band is twice the width of the standard deviation in the energy. Also plotted are the reference EDIP energies. (b) DUNN uncertainty. The uncertainty band is the same as that in panel (a) except that here it is centered around 0 instead of the prediction mean in panel (a).

across most of the range except for lattice parameters smaller than 5 Å and larger than 5.9 Å. The training set contains configurations with lattice parameters up to $\pm 10\%$ from the equilibrium value (i.e. 4.89 ~ 5.97 Å). Thus configurations with lattice parameters smaller than 5 Å and larger than 5.9 Å are at the “edge” of the training data where accuracy of the NN potential is clearly reduced. This is consistent with the discussion in Section 1. While highly accurate within the training set, the NN potential has low transferability and thus its ability to extrapolate beyond its training set is limited. This is particularly clear on the compressive end of the response (lattice constant smaller than 5.0 Å). In contrast, the SW potential has a lower accuracy overall since it is constrained by its physical functional form, but this leads to a more correct trend outside the training set.

It is important to quantify the uncertainty in the predictions of machine learning potentials given their low transferability. As discussed in Section 3.2, KLIF supports the training of DUNN potentials [88] that are based on dropout uncertainty estimation. To demonstrate this, we train a DUNN potential for the silicon dataset and apply it to investigate the same energy versus lattice parameter problem discussed above. Since the emphasis is on the uncertainty in energy, forces are not used in the training. (Details of the parameterization procedure are provided in the supplementary material.) When a DUNN model is used it provides a mean value, which is the average over the dropout ensemble, and an associated uncertainty estimate. The results for the cohesive energy versus lattice parameter are compared with the EDIP reference data in Fig. 8(a). The mean DUNN values are in excellent agreement with the reference data.¹¹ More importantly, the band around the mean values shows that the DUNN uncertainty estimate increases as the silicon crystal is strained away from its equilibrium state ($a = 5.43$ Å) and that the increase accelerates towards the edges of the training set (see Fig. 8(b)). Such uncertainty information can help to determine whether a molecular simulation is reliable or not.

¹¹ The agreement is better than the NN potential in Fig. 2 since only energies are used in training the DUNN potential allowing it to obtain a better fit, whereas the NN potential is fit using energies and forces.




Fig. 9. Phonon dispersions of diamond cubic silicon along high symmetry points in the first Brillouin zone predicted by the trained SW and NN potentials along with the reference data by EDIP. (For interpretation of the colors in the figure, the reader is referred to the web version of this article.)

As a second example, we consider phonon dispersion. This set of curves provides a comprehensive view of the elastic vibrational properties of a material, which play a key role in many dynamical properties including thermal transport and stress wave propagation. It is therefore important for IPs to predict phonon dispersion correctly. Fig. 9 presents the phonon dispersion curves of silicon along high-symmetry points in the first Brillouin zone obtained using the phonopy package [114]. The SW potential is in better agreement with the reference data for branches with larger phonon frequencies, but is less accurate for the two lowest-frequency branches, which can be seen at the W, K, and U points. Despite these small differences, the predictions by both the SW and NN potentials are in good agreement with the reference data. The training set does not explicitly contain phonon frequency data, so the fact that both the SW and NN potentials are able to correctly reproduce EDIP’s phonon dispersion curves indicates that they provide a good representation for the EDIP functional form near the equilibrium state.

6. Summary and outlook

In this paper, we introduce the *KIM-based learning-integrated fitting framework* (KLIF) for developing IPs. KLIF provides a uniform Python user interface to train both physics-based and machine learning potentials. It is flexible and easily extended to support new atomic environment descriptors, models, loss functions, minimizers, and analyzers. KLIF integrates closely with the KIM framework. An IP trained using KLIF can be readily deployed in a format consistent with the KIM API, which enables it to be used directly in major simulation codes such as LAMMPS [35,44,45], ASE [38,39], DL_POLY [40,41], GULP [42,43] and ASAP [115] among others. The package is distributed under an open-source license and is available at <https://github.com/openkim/kliff> along with a comprehensive user manual with several tutorials.

KLIF (version 0.3.0) is fully functional as demonstrated in this paper by training the SW, NN, and DUNN potentials for silicon. Development continues with an emphasis on incorporating new features, including (1) supporting more machine learning models and descriptors; (2) integration with KIM tests to train on material properties beyond energy, forces, and stress; and (3) creation of tools for automatic selection of hyperparameters for machine learning potentials (e.g. optimal number of terms to retain for a

descriptor and optimal number of layers and nodes in each layer for an NN potential). We encourage other researchers to contribute to the development, and provide full and detailed documentation of the KLIFF API (see the Package Reference section in the documentation <https://github.com/openkim/kliff>).

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This research was partly supported by the Army Research Office (W911NF-14-1-0247) under the MURI program, the National Science Foundation (NSF) under grants DMR-1834251, DMR-1834332 and OAC-2039575, and through the University of Minnesota MR-SEC under Award Number DMR-1420013. The authors wish to acknowledge the Minnesota Supercomputing Institute (MSI) at the University of Minnesota for providing resources that contributed to the results reported in this paper. MW thanks the University of Minnesota Doctoral Dissertation Fellowship for supporting his research.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.cpc.2021.108218>.

References

- [1] E.B. Tadmor, R.E. Miller, *Modeling Materials: Continuum, Atomistic and Multiscale Techniques*, Cambridge University Press, 2011.
- [2] R.Z. Khaliullin, H. Eshet, T.D. Kühne, J. Behler, M. Parrinello, *Nat. Mater.* 10 (9) (2011) 693, <https://doi.org/10.1038/nmat3078>.
- [3] S. Piana, K. Lindorff-Larsen, D.E. Shaw, *Proc. Natl. Acad. Sci.* 109 (44) (2012) 17845–17850, <https://doi.org/10.1073/pnas.1201811109>.
- [4] M. Wen, S.N. Shirodkar, P. Plecháč, E. Kaxiras, R.S. Elliott, E.B. Tadmor, *J. Appl. Phys.* 122 (24) (2017) 244301, <https://doi.org/10.1063/1.5007842>.
- [5] M. Wen, E.B. Tadmor, *Phys. Rev. B* 100 (19) (2019) 195419, <https://doi.org/10.1103/physrevb.100.195419>.
- [6] M. Born, R. Oppenheimer, *Ann. Phys.* 389 (20) (1927) 457–484, <https://doi.org/10.1002/andp.19273892002>.
- [7] Y. Mishin, D. Farkas, M.J. Mehl, D.A. Papaconstantopoulos, *Phys. Rev. B* 59 (1999) 3393–3407, <https://doi.org/10.1103/physrevb.59.3393>.
- [8] M. Wen, S.M. Whalen, R.S. Elliott, E.B. Tadmor, *Model. Simul. Mater. Sci. Eng.* 23 (7) (2015) 074008, <https://doi.org/10.1088/0965-0393/23/7/074008>.
- [9] M. Wen, J. Li, P. Brommer, R.S. Elliott, J.P. Sethna, E.B. Tadmor, *Model. Simul. Mater. Sci. Eng.* 25 (1) (2017) 014001, <https://doi.org/10.1088/0965-0393/25/1/014001>.
- [10] F. Ercolessi, J.B. Adams, *Europhys. Lett.* 26 (8) (1994) 583–588, <https://doi.org/10.1209/0295-5075/26/8/005>.
- [11] P. Zhang, D.R. Trinkle, *Model. Simul. Mater. Sci. Eng.* 23 (2015) 065011, <https://doi.org/10.1088/0965-0393/23/6/065011>.
- [12] J.E. Jones, *Proc. R. Soc. A* 106 (738) (1924) 441–462, <https://doi.org/10.1098/rspa.1924.0081>.
- [13] J.E. Jones, *Proc. R. Soc. A* 106 (738) (1924) 463–477, <https://doi.org/10.1098/rspa.1924.0082>.
- [14] J.E. Lennard-Jones, *Proc. Phys. Soc.* 43 (5) (1931) 461–482, <https://doi.org/10.1088/0959-5309/43/5/301>.
- [15] A.C.T. van Duin, S. Dasgupta, F. Lorant, W.A. Goddard, *J. Phys. Chem. A* 105 (41) (2001) 9396–9409, <https://doi.org/10.1021/jp004368u>.
- [16] D.W. Brenner, *Phys. Status Solidi (b)* 217 (2000) 23–40, <https://doi.org/10.1002/3527603107.ch2>.
- [17] J. Behler, M. Parrinello, *Phys. Rev. Lett.* 98 (14) (2007) 146401, <https://doi.org/10.1103/physrevlett.98.146401>.
- [18] A.P. Bartók, M.C. Payne, R. Kondor, G. Csányi, *Phys. Rev. Lett.* 104 (13) (2010) 136403, <https://doi.org/10.1103/physrevlett.104.136403>.
- [19] M. Rupp, A. Tkatchenko, K.-R. Müller, O.A. Von Lilienfeld, *Phys. Rev. Lett.* 108 (5) (2012) 058301, <https://doi.org/10.1103/physrevlett.108.058301>.
- [20] A.P. Thompson, L.P. Swiler, C.R. Trott, S.M. Foiles, G.J. Tucker, *J. Comput. Phys.* 285 (2015) 316–330, <https://doi.org/10.1016/j.jcp.2014.12.018>.
- [21] A.V. Shapeev, *Multiscale Model. Simul.* 14 (3) (2016) 1153–1173, <https://doi.org/10.1137/15m1054183>.
- [22] S. Hajinazar, J. Shao, A.N. Kolmogorov, *Phys. Rev. B* 95 (1) (2017) 014114, <https://doi.org/10.1103/physrevb.95.014114>.
- [23] J.J. Waterfall, F.P. Casey, R.N. Gutenkunst, K.S. Brown, C.R. Myers, P.W. Brouwer, V. Elser, J.P. Sethna, *Phys. Rev. Lett.* 97 (2006) 150601, <https://doi.org/10.1103/PhysRevLett.97.150601>.
- [24] Y. Kurniawan, C.L. Petrie, K.J. Williams, M.K. Transtrum, E.B. Tadmor, R.S. Elliott, D.S. Karls, M. Wen, Bayesian, frequentist, and information geometry approaches to parametric uncertainty quantification of classical empirical potentials, submitted for publication, 2021.
- [25] M.K. Transtrum, B.B. Machta, J.P. Sethna, *Phys. Rev. E* 83 (3) (2011) 036701, <https://doi.org/10.1103/PhysRevE.83.036701>.
- [26] M.K. Transtrum, J.P. Sethna, Geodesic acceleration and the small-curvature approximation for nonlinear least squares, arXiv preprint, arXiv:1207.4999, 2012.
- [27] M.K. Transtrum, J.P. Sethna, Improvements to the Levenberg–Marquardt algorithm for nonlinear least-squares minimization, arXiv preprint, arXiv:1201.5885, 2012.
- [28] *The Minerals, Metals & Materials Society (TMS), Verification & Validation of Computational Models Associated with the Mechanics of Materials*, TMS, Pittsburgh, PA, 2019.
- [29] E.B. Tadmor, R.S. Elliott, J.P. Sethna, R.E. Miller, C.A. Becker, *JOM* 63 (7) (2011) 17, <https://doi.org/10.1007/s11837-011-0102-6>.
- [30] E.B. Tadmor, R.S. Elliott, S.R. Phillpot, S.B. Sinnott, *Curr. Opin. Solid State Mater. Sci.* 17 (6) (2013) 298–304, <https://doi.org/10.1016/j.cossms.2013.10.004>.
- [31] D.S. Karls, M. Bierbaum, A.A. Alemi, R.S. Elliott, J.P. Sethna, E.B. Tadmor, *J. Chem. Phys.* 153 (2020) 064104, <https://doi.org/10.1063/5.0014267>.
- [32] Open knowledgebase of interatomic models (OpenKIM), <https://openkim.org>, 2021. (Accessed 30 May 2021).
- [33] C. de Tomas, I. Suarez-Martinez, N.A. Marks, *Carbon* 109 (2016) 681–693, <https://doi.org/10.1016/j.carbon.2016.08.024>.
- [34] D.W. Brenner, O.A. Shenderova, J.A. Harrison, S.J. Stuart, B. Ni, S.B. Sinnott, *J. Phys. Condens. Matter* 14 (4) (2002) 783–802, <https://doi.org/10.1088/0953-8984/14/4/312>.
- [35] Large-scale atomic/molecular massively parallel simulator (LAMMPS), <https://www.lammps.org>, 2021. (Accessed 30 May 2021).
- [36] S.J. Stuart, A.B. Tutein, J.A. Harrison, *J. Chem. Phys.* 112 (14) (2000) 6472–6486, <https://doi.org/10.1063/1.481208>.
- [37] R.S. Elliott, E.B. Tadmor, Knowledgebase of Interatomic Models (KIM) application programming interface (API), <https://doi.org/10.25950/ff8f563a>, 2011.
- [38] A.H. Larsen, J.J. Mortensen, J. Blomqvist, I.E. Castelli, R. Christensen, M. Dulak, J. Friis, M.N. Groves, B. Hammer, C. Hargus, E.D. Hermes, P.C. Jennings, P.B. Jensen, J. Kermode, J.R. Kitchin, E.L. Kolsbjerg, J. Kubal, K. Kaasbjerg, S. Lysgaard, J.B. Maronsson, T. Maxson, T. Olsen, L. Pastewka, A. Peterson, C. Rossgaard, J. Schiøtz, O. Schütt, M. Strange, K.S. Thygesen, T. Vegge, L. Vilhelmsen, M. Walter, Z. Zeng, K.W. Jacobsen, *J. Phys. Condens. Matter* 29 (27) (2017) 273002, <https://doi.org/10.1088/1361-648x/aa680e>.
- [39] ASE: the atomic simulation environment—a Python library for working with atoms, <https://wiki.fysik.dtu.dk/ase/>, 2021. (Accessed 30 May 2021).
- [40] W. Smith, T.R. Forester, *J. Mol. Graph.* 14 (1996) 136–141, [https://doi.org/10.1016/S0263-7855\(96\)00043-4](https://doi.org/10.1016/S0263-7855(96)00043-4).
- [41] DL_POLY classic molecular simulation package, https://www.scd.stfc.ac.uk/Pages/DL_POLY.aspx, 2021. (Accessed 30 May 2021).
- [42] J.D. Gale, *J. Chem. Soc. Faraday Trans. 93* (4) (1997) 629–637, <https://doi.org/10.1039/a606455h>.
- [43] General utility lattice program (GULP), <https://gulp.curtin.edu.au/gulp>, 2021. (Accessed 30 May 2021).
- [44] S. Plimpton, *J. Comput. Phys.* 117 (1) (1995) 1–19, <https://doi.org/10.1006/jcph.1995.1039>.
- [45] A.P. Thompson, H.M. Aktulga, R. Berger, D.S. Bolintineanu, W. Michael Brown, P.S. Crozier, P.J. in 't Veld, A. Kohlmeyer, S.G. Moore, T.D. Nguyen, R. Shan, M. Stevens, J. Tranchida, C. Trott, S.J. Plimpton, *Comput. Phys. Commun.* (2021) 108171, <https://doi.org/10.1016/j.cpc.2021.108171>.
- [46] E.B. Tadmor, M. Ortiz, R. Phillips, *Philos. Mag.* A 73 (6) (1996) 1529–1563, <https://doi.org/10.1080/01418619608243000>.
- [47] Quasicontinuum method website, <https://openkim.org>, 2009.
- [48] F.H. Stillinger, T.A. Weber, *Phys. Rev. B* 31 (8) (1985) 5262, <https://doi.org/10.1103/physrevb.31.5262>.
- [49] M. Wen, Stillinger–Weber (SW) Model Driver v005, Online, <https://doi.org/10.25950/934dca3e>, 2018. (Accessed 30 May 2021).
- [50] A.P. Bartók, R. Kondor, G. Csányi, *Phys. Rev. B* 87 (18) (2013) 184115, <https://doi.org/10.1103/physrevb.87.184115>.
- [51] J. Behler, *J. Chem. Phys.* 134 (7) (2011) 074106, <https://doi.org/10.1063/1.3553717>.
- [52] H. Huo, M. Rupp, Unified representation of molecules and crystals for machine learning, arXiv preprint, arXiv:1704.06439, 2017.
- [53] M.F. Langer, A. Goelmann, M. Rupp, Representations of molecules and materials for interpolation of quantum-mechanical simulations via machine learning, arXiv preprint, arXiv:2003.12081, 2020.

- [54] M. Wen, Development of interatomic potentials with uncertainty quantification: applications to two-dimensional materials, Ph.D. thesis, University of Minnesota, 2019, <https://hdl.handle.net/11299/206694>.
- [55] L. Prechelt, *Neural Netw.* 11 (4) (1998) 761–767, [https://doi.org/10.1016/S0893-6080\(98\)00010-0](https://doi.org/10.1016/S0893-6080(98)00010-0).
- [56] G.E. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever, R.R. Salakhutdinov, Improving neural networks by preventing co-adaptation of feature detectors, *arXiv preprint, arXiv:1207.0580*, 2012.
- [57] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, *J. Mach. Learn. Res.* 15 (1) (2014) 1929–1958.
- [58] P. Brommer, F. Gähler, *Model. Simul. Mater. Sci. Eng.* 15 (3) (2007) 295, <https://doi.org/10.1088/0965-0393/15/3/008>.
- [59] P. Brommer, A. Kiselev, D. Schopf, P. Beck, J. Roth, H.-R. Trebin, *Model. Simul. Mater. Sci. Eng.* 23 (7) (2015) 074002, <https://doi.org/10.1088/0965-0393/23/7/074002>.
- [60] N. Artrith, *A. Urban, Comput. Mater. Sci.* 114 (2016) 135–150, <https://doi.org/10.1016/j.commatsci.2015.11.047>.
- [61] A. Khorshidi, A.A. Peterson, *Comput. Phys. Commun.* 207 (2016) 310–324, <https://doi.org/10.1016/j.cpc.2016.05.010>.
- [62] A.E. Allen, G. Dussan, C. Ortner, G. Csányi, *Mach. Learn. Sci. Technol.* 2 (2) (2021) 025017, <https://doi.org/10.1088/2632-2153/abd51e>.
- [63] A. Stukowski, E. Fransson, M. Mock, P. Erhart, *Model. Simul. Mater. Sci. Eng.* 25 (5) (2017) 055003, <https://doi.org/10.1088/1361-651x/aa6ecf>.
- [64] H. Wang, L. Zhang, J. Han, E. Weinan, *Comput. Phys. Commun.* 228 (2018) 178–184, <https://doi.org/10.1016/j.cpc.2018.03.016>.
- [65] GAP and SOAP documentation, <https://libatoms.github.io/GAP/>, 2021. (Accessed 16 September 2021).
- [66] S. Hajinazar, A. Thorn, E.D. Sandoval, S. Kharabadze, A.N. Kolmogorov, *Comput. Phys. Commun.* 259 (2021) 107679, <https://doi.org/10.1016/j.cpc.2020.107679>.
- [67] I.S. Novikov, K. Gubaev, E.V. Podryabinkin, A.V. Shapeev, *Mach. Learn. Sci. Technol.* 2 (2) (2020) 025002, <https://doi.org/10.1088/2632-2153/abc9fe>.
- [68] Y. Lysogorskiy, C. van der Oord, A. Bochkarev, S. Menon, M. Rinaldi, T. Hammerschmidt, M. Mrovec, A. Thompson, G. Csányi, C. Ortner, et al., *npj Comput. Mater.* 7 (97) (2021), <https://doi.org/10.5281/zenodo.4734036>.
- [69] R. Lot, F. Pellegrini, Y. Shaidu, E. Küçükbenli, *Comput. Phys. Commun.* 256 (2020) 107402, <https://doi.org/10.1016/j.cpc.2020.107402>.
- [70] H. Yanxon, D. Zagaceta, B. Tang, D.S. Matteson, Q. Zhu, *Mach. Learn. Sci. Technol.* 2 (2) (2020) 027001, <https://doi.org/10.1088/2632-2153/abc940>.
- [71] The RuNNer code, <https://www.uni-goettingen.de/de/560580.html>, 2021. (Accessed 16 September 2021).
- [72] K. Lee, D. Yoo, W. Jeong, S. Han, *Comput. Phys. Commun.* 242 (2019) 95–103, <https://doi.org/10.1016/j.cpc.2019.04.014>.
- [73] E. Tadmor, Verification Check for Memory Leaks using Valgrind v001, Online, <https://doi.org/10.25950/ba474f45>, 2018. (Accessed 30 May 2021).
- [74] E. Tadmor, Verification Check of Invariance with respect to the Inversion Operation (Inversion Symmetry) v001, Online, <https://doi.org/10.25950/63a96579>, 2018. (Accessed 30 May 2021).
- [75] E. Tadmor, Verification Check of Invariance with respect to Atom Permutations (Permutation Symmetry) v001, Online, <https://doi.org/10.25950/dfbf8222>, 2018. (Accessed 30 May 2021).
- [76] E. Tadmor, Verification Check of Forces via Numerical Differentiation (Richardson Extrapolation Technique) v002, Online, <https://doi.org/10.25950/9be59b8d>, 2018. (Accessed 30 May 2021).
- [77] S. Pattamatta, Stacking and twinning fault energies of an fcc lattice at zero temperature and pressure v001, Online, <https://doi.org/10.25950/d6ffade7>, 2018. (Accessed 30 May 2021).
- [78] J. Li, E. Tadmor, Elastic constants for cubic crystals at zero temperature and pressure v005, Online, <https://doi.org/10.25950/49c5c255>, 2019. (Accessed 30 May 2021).
- [79] M. Wen, Linear thermal expansion coefficient of a cubic crystal structure at a given temperature and pressure v001, Online, <https://doi.org/10.25950/fc69d82d>, 2016. (Accessed 30 May 2021).
- [80] R.A. Messerly, T.A. Knotts, W.V. Wilding, *J. Chem. Phys.* 146 (19) (2017) 194110, <https://doi.org/10.1063/1.4983406>.
- [81] H. Cramér, *Mathematical Methods of Statistics*, Princeton University Press, Princeton, 1999.
- [82] S.L. Frederiksen, K.W. Jacobsen, K.S. Brown, J.P. Sethna, *Phys. Rev. Lett.* 93 (16) (2004) 165501, <https://doi.org/10.1103/PhysRevLett.93.165501>.
- [83] S. Longbottom, P. Brommer, *Model. Simul. Mater. Sci. Eng.* 27 (4) (2019) 044001, <https://doi.org/10.1088/1361-651x/ab0d75>.
- [84] N. Artrith, J. Behler, *Phys. Rev. B* 85 (4) (2012) 045439, <https://doi.org/10.1103/PhysRevB.85.045439>.
- [85] A.A. Peterson, R. Christensen, A. Khorshidi, *Phys. Chem. Chem. Phys.* 19 (18) (2017) 10978–10985, <https://doi.org/10.1039/C7CP00375G>.
- [86] L. Zhang, D.-Y. Lin, H. Wang, R. Car, E. Weinan, *Phys. Rev. Mater.* 3 (2) (2019) 023804.
- [87] W. Jeong, D. Yoo, K. Lee, J. Jung, S. Han, *J. Phys. Chem. Lett.* 11 (15) (2020) 6090–6096, <https://doi.org/10.1021/acs.jpclett.0c01614>.
- [88] M. Wen, E.B. Tadmor, *npj Comput. Mater.* 6 (1) (2020) 124, <https://doi.org/10.1038/s41524-020-00390-8>.
- [89] Y. Gal, Z. Ghahramani, in: M.F. Balcan, K.Q. Weinberger (Eds.), *Proceedings of the 33rd International Conference on Machine Learning*, in: *Proceedings of Machine Learning Research*, vol. 48, PMLR, New York, New York, USA, 2016, pp. 1050–1059, <https://proceedings.mlr.press/v48/gal16.html>.
- [90] Y. Gal, *Uncertainty in deep learning*, Ph.D. thesis, University of Cambridge, 2016.
- [91] M. Wen, A dropout uncertainty neural network (DUNN) model driver v000, OpenKIM, <https://doi.org/10.25950/9573ca43>, 2019.
- [92] T. Brink, Model driver for Tersoff-style potentials ported from LAMMPS v003, Online, <https://doi.org/10.25950/55b7b34e>, 2019. (Accessed 30 May 2021).
- [93] J. Tersoff, *Phys. Rev. B* 37 (12) (1988) 6991–7000, <https://doi.org/10.1103/PhysRevB.37.6991>.
- [94] J. Tersoff, *Phys. Rev. B* 39 (1989) 5566–5568, <https://doi.org/10.1103/PhysRevB.39.5566>.
- [95] J. Nord, K. Albe, P. Erhart, K. Nordlund, *J. Phys. Condens. Matter* 15 (2003) 5649, <https://doi.org/10.1088/0953-8984/15/32/324>.
- [96] D.S. Karls, Environment-Dependent Interatomic Potential (EDIP) model driver v002, Online, <https://doi.org/10.25950/75c4686e>, 2018. (Accessed 30 May 2021).
- [97] M.Z. Bazant, E. Kaxiras, *Phys. Rev. Lett.* 77 (1996) 4370–4373, <https://doi.org/10.1103/PhysRevLett.77.4370>.
- [98] M.Z. Bazant, E. Kaxiras, J.F. Justo, *Phys. Rev. B* 56 (1997) 8542–8552, <https://doi.org/10.1103/PhysRevB.56.8542>.
- [99] J.A.F. Justo, M.Z. Bazant, E. Kaxiras, V.V. Bulatov, S. Yip, *Phys. Rev. B* 58 (1998) 2539–2550, <https://doi.org/10.1103/PhysRevB.58.2539>.
- [100] R.S. Elliott, EAM Model Driver for tabulated potentials with cubic Hermite spline interpolation as used in LAMMPS v005, Online, <https://doi.org/10.25950/68defa36>, 2018. (Accessed 30 May 2021).
- [101] M.S. Daw, M.I. Baskes, *Phys. Rev. B* 29 (12) (1984) 6443, <https://doi.org/10.1103/physrevb.29.6443>.
- [102] M.S. Daw, S.M. Foiles, M.I. Baskes, *Mater. Sci. Rep.* 9 (7) (1993) 251–310, [https://doi.org/10.1016/0920-2307\(93\)90001-u](https://doi.org/10.1016/0920-2307(93)90001-u).
- [103] L. Himanen, M.O. Jäger, E.V. Morooka, F.F. Canova, Y.S. Ranawat, D.Z. Gao, P. Rinke, A.S. Foster, *Comput. Phys. Commun.* 247 (2020) 106949, <https://doi.org/10.1016/j.cpc.2019.106949>.
- [104] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimeshain, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, S. Chintala, in: *Advances in Neural Information Processing Systems*, vol. 32, Curran Associates, Inc., 2019, pp. 8024–8035.
- [105] C. Zhu, R.H. Byrd, P. Lu, J. Nocedal, *ACM Trans. Math. Softw.* 23 (4) (1997) 550–560, <https://doi.org/10.1145/279232.279236>.
- [106] SciPy: a Python-based ecosystem of open-source software for mathematics, science, and engineering, <https://www.scipy.org>, 2021. (Accessed 30 May 2021).
- [107] H. Robbins, S. Monro, *Ann. Math. Stat.* (1951) 400–407, https://doi.org/10.1007/978-1-4612-5110-1_9.
- [108] J. Kiefer, J. Wolfowitz, et al., *Ann. Math. Stat.* 23 (3) (1952) 462–466, https://doi.org/10.1007/978-1-4613-8505-9_4.
- [109] D.P. Kingma, J. Ba, Adam: a method for stochastic optimization, *arXiv preprint, arXiv:1412.6980*, 2014.
- [110] K. He, X. Zhang, S. Ren, J. Sun, in: *2015 IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [111] M. Wen, S. Carr, S. Fang, E. Kaxiras, E.B. Tadmor, *Phys. Rev. B* 98 (23) (2018) 235404, <https://doi.org/10.1103/physrevb.98.235404>.
- [112] X. Zhou, D. Ward, J. Martin, F. Van Swol, J. Cruz-Campa, D. Zubia, *Phys. Rev. B* 88 (8) (2013) 085309, <https://doi.org/10.1103/physrevb.88.085309>.
- [113] M. Li, T. Zhang, Y. Chen, A.J. Smola, in: *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, 2014, pp. 661–670.
- [114] A. Togo, I. Tanaka, *Scr. Mater.* 108 (2015) 1–5, <https://doi.org/10.1016/j.scriptamat.2015.07.021>.
- [115] Asap: a calculator for doing large-scale classical molecular dynamics, <https://wiki.fysik.dtu.dk/asap/>, 2021. (Accessed 30 May 2021).