

Symbolic Regression in Materials Science

Burlacu, Bogdan; Kommenda, Michael; Kronberger, Gabriel; Winkler, Stephan; Affenzeller, Michael

DOI:

[10.48550/arXiv.2206.06422](https://doi.org/10.48550/arXiv.2206.06422)

Published: 13/06/2022

[Link to publication in pure](#)

Citation for published version (APA):

Burlacu, B., Kommenda, M., Kronberger, G., Winkler, S., & Affenzeller, M. (2022). *Symbolic Regression in Materials Science: Discovering Interatomic Potentials from Data*. arXiv. <https://doi.org/10.48550/arXiv.2206.06422>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain.
- You may freely distribute the URL identifying the publication in the public portal.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Symbolic Regression in Materials Science: Discovering Interatomic Potentials from Data

Bogdan Burlacu, Michael Kommenda, Gabriel Kronberger, Stephan Winkler,
Michael Affenzeller

Abstract Particle-based modeling of materials at atomic scale plays an important role in the development of new materials and understanding of their properties. The accuracy of particle simulations is determined by *interatomic potentials*, which allow to calculate the potential energy of an atomic system as a function of atomic coordinates and potentially other properties. First-principles-based *ab initio* potentials can reach arbitrary levels of accuracy, however their applicability is limited by their high computational cost.

Machine learning (ML) has recently emerged as an effective way to offset the high computational costs of *ab initio* atomic potentials by replacing expensive models with highly efficient surrogates trained on electronic structure data. Among a plethora of current methods, symbolic regression (SR) is gaining traction as a powerful “white-box” approach for discovering functional forms of interatomic potentials.

This contribution discusses the role of symbolic regression in Materials Science (MS) and offers a comprehensive overview of current methodological challenges and state-of-the-art results. A genetic programming-based approach for modeling atomic potentials from raw data (consisting of snapshots of atomic positions and associated potential energy) is presented and empirically validated on *ab initio* electronic structure data.

Heuristic and Evolutionary Algorithms Laboratory
University of Applied Sciences Upper Austria
Softwarepark 11, 4232 Hagenberg, Austria

1 Introduction

Materials Science (MS) is a highly interdisciplinary field incorporating elements of physics, chemistry, engineering and more recently, machine learning, in order to design and discover new materials. The rapid increase in processing power over the last decades has made computational modeling and simulation the main tool for studying new materials and determining their properties and behavior. Computational approaches can deliver accurate quantitative results without the need to set up and execute highly complex and costly physical experiments.

Potential energy surfaces (PES), describing the relationship between an atomic system's potential energy and the geometry of its atoms, are a central concept in computational chemistry and play a pivotal role in particle simulations. An example PES for the water molecule is shown in Figure 1. The mathematical function used to calculate the potential energy of a system of atoms with given positions in space and generate the PES is called an *interatomic potential* function. The form of this function, its physical fidelity as well as its complexity and efficiency are critical components in simulations used to predict material properties.

The ability to simulate large particle systems over long timescales depends critically on the accuracy and computational efficiency of the interatomic potential. Broadly speaking, the more accurate the methods, the lower its computational efficiency and the more limited its applicability. For example, first-principle modeling methods such as *density functional theory* (DFT) [33] provide highly accurate results by considering quantum-chemical effects but are not efficient enough to simulate large systems containing thousands of atoms over long timescales of nanoseconds [43].

Molecular dynamics (MD) simulations treat materials as systems consisting of many microscopic particles (atoms) which interact with each other through interatomic potentials depending mainly on their positions and are governed by the laws of statistical thermodynamics. Macroscopic properties of materials are obtained as time and/or ensemble averages of processes emerging at the microscopic scale [27].

Empirical and semi-empirical methods treat atomic interactions in a more coarse-grained manner via parameterized analytical functional forms and trade-off accuracy for execution speed in order to enable simulations at a larger scale. Although they are computationally undemanding, they are only able to provide a qualitatively reasonable description of chemical interactions [52].

Machine learning (ML) interatomic potentials aim to bridge the gap between quantum and empirical methods in order to deliver the best of both worlds: functional forms that are as efficient as empirical potentials and as accurate as quantum-chemical approaches.

1.1 Materials informatics and data-driven potentials

Building upon the three established paradigms of science that have led to many technological advances over time: experimental, theoretical and simulation-based, a

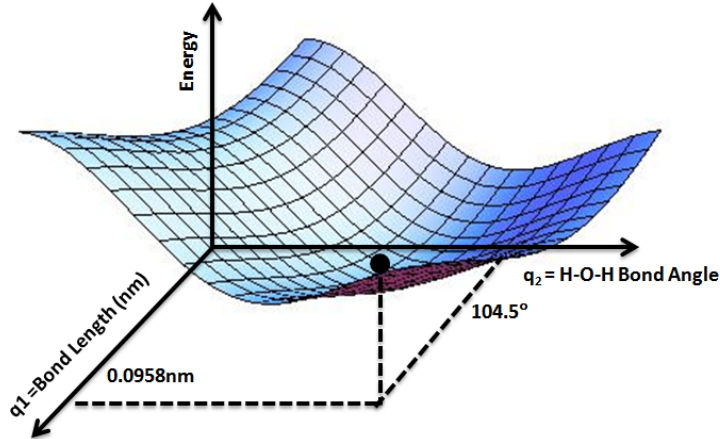


Fig. 1: PES for water molecule: Shows the energy minimum corresponding to optimized molecular structure for water- O-H bond length of 0.0958nm and H-O-H bond angle of 104.5°. Image from Wikipedia ©AimNature

fourth “data-driven“ paradigm of science is emerging today using machine learning and the large amounts of experimental and simulation-data available [1]. “Big-data” science unifies the first three paradigms and opens up new avenues in materials science under the umbrella term of *materials informatics*. The field of material informatics is very new and many unsolved questions still remain open and wait for proper answers [26].

Machine learning interaction models are generated on the basis of quantum-chemical reference data consisting of a series of snapshots of atomic coordinates, associated potential energy of the system and optionally other properties.

In molecular dynamics simulations, the system’s potential energy is typically decomposed into a set of independent m -body interactions that are a function of each particle’s position, \mathbf{r} . For a two-body or pair potential, it is assumed that the energy contributions from each pair of interacting particles are independent of other pairs and therefore:

$$E = \sum_{\langle i,j \rangle} g(\mathbf{r}_i, \mathbf{r}_j) \quad (1)$$

For a three-body potential, triplets of atoms are also considered:

$$E = \sum_{\langle i,j \rangle} g(\mathbf{r}_i, \mathbf{r}_j) + \sum_{\langle i,j,k \rangle} h(\mathbf{r}_i, \mathbf{r}_j, \mathbf{r}_k) \quad (2)$$

Traditionally, the functions g and h are represented by all kinds of empirical or semi-empirical analytical functions. With the advent of machine learning and data-based modeling, it becomes possible to automatically search for these functional forms with the help of ab initio training data. Substantial effort has already been

put into this direction and many machine learning models have been successful in discovering interatomic potentials for a variety of chemical configurations [41].

1.2 Current challenges

Despite their success in representing atomic interactions, ML-methods are not without their own challenges. Deriving highly-accurate and tractable analytic functional forms for high-dimensional PESs is a very active field of research. The most important requirements for ML-based PESs are:

- general applicability and absence of ad-hoc approximations (transferability)
- accuracy close to first-principles methods (including high-order many-body effects)
- very high efficiency to enable large simulations
- the ability to describe chemical reactions and arbitrary atomic configurations
- the ability to be automatically constructed and systematically improved

Currently available potentials are far from satisfying all the needs [6], mainly due to the following difficulties and shortcomings:

Physical plausibility

Closed physical systems are governed by various conservation laws that describe invariant properties. These fundamental principles of nature provide strong constraints that can be used to guide the search towards physically-plausible ML models [52]. In molecular systems each conserved quantity is associated with a differentiable symmetry of the action of a physical system.

Typical conserved quantities include temporal and roto-translational invariance (i.e. total energy, linear and angular momentum). Forces must be the negative gradient of the potential energy E with respect to atomic positions r_i :

$$F_i = -\nabla_{r_i} E$$

When atoms move, they always acquire the same amount of kinetic energy as they lose in potential energy, and vice versa – the total energy is conserved. The potential energy of a molecule only depends on the relative positions of atoms and does not change with rigid rotations or translations.

Another aspect of invariance is permutational invariance resulting from the fact that from the perspective of the electrons, atoms with the same nuclear charge appear identical to each other and can thus be exchanged without affecting the energy or the forces. To ensure physically meaningful predictions, ML-based models must exhibit the same invariant behavior as the true potential energy surface.

Accuracy

Accuracy is one of the most important requirements of ML potentials. The predicted energies and forces should be as close as possible to the underlying *ab initio* data. Numerical accuracy of the ML models is restricted by the intrinsic limitations of their functional form and descriptors (input variables) used. For example, conceptual problems related to incorporating rotational, translational and permutational invariance into descriptors are of primary relevance [6, 21, 45, 46] as well as their optimal design [20].

Transferability

Ideally, potentials should be generally applicable and should not be restricted to specific types of atomic configurations. Due to their mathematical unbiased form, ML methods are promising candidates to reach this goal. However in practice, developed potentials often perform very well in applications they have been designed for, but are too system-specific and thus cannot be easily transferred from one system to another. The issues of extensibility, generality and transferability of the ML potentials need to be explicitly addressed [6].

Complexity and data requirements

Another issue worth mentioning here is the mathematical complexity of ML potentials. For example, the most popular ML methods used to represent many-body PESs, ANNs, require complex architectures with many adjustable parameters (weights of neural synapses and neuron biases) to yield sufficiently flexible and invariant PESs representations. For this, large amounts of training data (often dozens or even hundreds of thousands points) are needed. On the other hand, the number of training data should be kept as low as possible since they are calculated via demanding quantum-chemical methods. It means that as simple as possible analytic representations of PESs are needed.

Integration of physical knowledge and interpretability

Related to the mathematical complexity issue, it is also important to note that most of the ML methods (e.g., ANN, SVM) are of a “black-box” nature and may be less amenable to including physical information into the functional forms, relying at least partially on physics-inspired features considered in atomic descriptors. This often leads to increased mathematical and computational complexity of resulting interaction models. One of the main directions of the current development in ML-based computational MS is the shift from “black-box” methods towards “white-box” methods which often offer better interpretability.

2 State of the art

A plethora of machine learning approaches have recently emerged as a powerful alternative for finding a functional relation between an atomic configuration and corresponding energy [6, 17, 23].

Several ML techniques such as polynomial fitting [10], Gaussian processes [5], spectral neighbor analysis [51], modified Shepard interpolation [29], moment tensor potentials [46], interpolating moving least squares [32], support vector machines [4], random forests [31], artificial neural networks (ANNs) [15, 25, 45, 54] or symbolic regression (SR) [39] have been successfully employed for a variety of systems.

More detailed reviews of current ML potentials can be found for example in [23], [36], [41] or [52]. Particularly, ANNs have received a considerable attention and are probably the most popular form of ML potentials used in MS [54]. However, methods based on symbolic regression are gaining in popularity due to the advantages they bring in solving aspects of physical knowledge integration, efficiency and interpretability [7, 8, 11, 12, 24, 37, 40, 44, 47].

In the following, we refer to symbolic regression in its canonical incarnation that employs genetic programming to perform a search over the space of mathematical expressions. Symbolic regression approaches have succeeded in rediscovering simple forms of potentials that deliver qualitatively good results in a series of specific applications, some of which are described below.

2.1 Directed search

The goal of directed search is to improve search efficiency by limiting the hypothesis space to a functional form known to deliver qualitatively good results, instead of searching for a brand new potential.

Makarov and Metiu [37] use the Morse potential as a functional template for modeling diatomic molecules (see Section 5, Eq. 17). They rewrite it in the form:

$$M(D(r), R(r)) = D(r)(1 - \exp(R(r)))^2 \quad (3)$$

and use genetic programming to find the best $D(r)$ and $R(r)$.

The directed search approach is augmented with an error metric that better reflects the physical characteristics of the problem. A standard error metric such as the MSE has the disadvantage of overemphasizing high-energy points which are rarely used during simulation. For this reason, the authors found it advantageous to introduce a scaling factor:

$$F(a) = \sum_i \frac{(E(r_i) - f(r_i; a))^2}{E^2(r_i) + \delta^2} \quad (4)$$

where the constant δ is added to prevent division by zero.

For each function f_α in the population of individuals, the fitness function is then defined as:

$$p_\alpha = \exp(-\beta F_\alpha) \quad (5)$$

where parameter β controls how discriminating the function is and is adaptively updated during the run. The search starts with a small value for β which is gradually increased as the search improves.

The authors note the importance of including the derivative of the energy in the training data:

$$F'(a) = \sum_i \frac{|\nabla E(r_i) - \nabla f(r_i; a)|^2}{|\nabla E(r_i)|^2 + \delta'^2} \quad (6)$$

Leading to an expanded fitness function p_α :

$$p_\alpha = \exp(-\beta(F + F')) \quad (7)$$

The recombination pool is filled using a proportional selection scheme. An additional ‘‘natural selection’’ operator employs a ‘‘badness list’’ $b_\alpha = \exp(\beta F_\alpha)$ whose elements are the inverse of the fitness. Old individuals are replaced with a probability proportional with badness.

Results

The directed search approach is shown to perform better than an undirected search over the search space, on training data generated using the Lippincott potential (Section 5, Eq. 19). A population size of 500 individuals is evolved over 150 generations (75,000 evaluations) using the primitive set $\mathcal{P} = \{+, -, \div, \times, \exp\}$. Furthermore, a search directed by a Lennard-Jones potential gives accuracy comparable to that directed by a Morse function, suggesting that restricting the hypothesis space with an appropriate functional template is a powerful and general approach in the search for interatomic potentials. In the case of the Lennard-Jones potential (Section 5, Equation 18) the functional template was defined as

$$f(r) = 4D(r) \left[\frac{1}{4} + \left(\frac{1}{R(r)} \right)^{12} - \left(\frac{1}{R(r)} \right)^6 \right] \quad (8)$$

The authors additionally note that some of the returned models, although accurate, exhibited unphysical behavior and did not extrapolate well. For example, one of the returned models based on the Lennard-Jones functional form had very good accuracy but contained a singularity at $r = 12 \text{ \AA}$, a point outside the interpolation range. The authors address overfitting by fitting the parameters of both the energy function and its derivative in the local search phase. This reduces the chance of obtaining pathological curves in the model extrapolation response.

Finally, Makarov and Metiu also model the potential of a triatomic molecule on *ab initio* data consisting of 60 nuclear configurations, showing that directed search maintains high levels of accuracy and scales favorably with dimensionality.

2.2 Directed Search with Parallel Multilevel Genetic Program

Belluci and Coker [7, 8] employ symbolic regression to discover empirical valence bond (EVB) models using directed search augmented with a multilevel genetic programming approach: the lower level (LLGP) optimizes co-evolving populations of models, while the higher level (HLGP) optimizes genetic operator probabilities of the lower level populations. The approach entitled Parallel Multilevel Genetic Program (PMLGP) found accurate EVB models for proton transfer in 3-hydroxygramma-pyrone (3-HGP) in gas phase and protic solvent as well as ultrafast enol keto isomerisation in the lowest singlet excited state of 3-hydroxyflavone (3-HF).

At the lower level (LLGP), the authors use the same error metric and fitness as in [37], namely Equations 4 and 5. LLGP individuals represent the $R(r)$ functional part of the Morse potential (see Equation 3). Remarkably, PMLGP does not use crossover but instead uses six different mutation operators:

- *Point mutation* randomly replaces a subtree with a randomly-generated one.
- *Branch mutation* replaces a binary operator with one of its arguments at random.
- *Leaf mutation* replaces a leaf node with another randomly selected leaf.
- *New tree mutation* replaces an entire tree with a newly generated tree.
- *Parameter change* replaces each parameter value a_i with $a_i + (R - 0.5)\gamma$, where R is a uniform random number on the unit interval and γ is a scaling constant.
- *Parameter scaling* replaces each parameter value a_i with $a_i R\gamma$, where R is a uniform random number on the unit interval and γ is a scaling constant.

Of the last two types of mutation, parameter change is designed to make small local moves in parameter space, while parameter scaling is designed to make large moves in parameter space as to escape the basins of attraction of local optima. Selection is performed using stochastic universal sampling [3].

At the higher level (HLGP) a real vector encoding is used to represent genetic operator probabilities. The population is initialized with k random vectors $P_k = (p_1^{(k)}, \dots, p_6^{(k)})$, with $\sum_i p_i^{(k)} = 1$, where k ranges from 1 to the total number N_p of processors, such that each vector corresponds to one of the LLGP populations whose operator probabilities it dynamically adapts.

The fitness of each vector P_k is evaluated based on the maximum fitness delta in the corresponding LLGP population over a specified time interval Δt :

$$F_k^{\text{HLGP}} = \frac{\Delta F_{\text{max}}^{\text{LLGP}}}{\Delta t} \quad (9)$$

This is based on the idea that the larger the magnitude of F_k^{HLGP} , the more successful the set of probabilities P_k at improving the fitness of the population.

Two genetic operators are used to modify the probability vectors P_k :

- *Mutation* changes each component of the vector by a random amount with the constraint that all components sum up to one. This operator kicks in when the fitness of a vector P_k drops below a given threshold.
- *Adaptation* attempts to improve the probability distribution given by P_k by using feedback from the LLGP. Each LLGP builds a histogram of the number of times each mutation produced the most fit member of the population. Then the success frequency of the mutation operator is given by:

$$s_i = \frac{w_i m_i}{n}, \quad w_i = \frac{1}{p_i}, \quad n = \sum_i m_i$$

Here, w_i is a weight, m_i is the number of successful mutations for the i^{th} operator (component of P_k) and n is the total number of successful mutations (for all operators). Based on the success frequencies, adaptation shifts a random amount of probability from the least successful operator to the most successful operator.

The number of LLGP populations (and HLGP individuals, respectively) is set to the number of available processors. Initially, all LLGP populations are identical but diverge during evolution as each corresponding fitness function is parameterized with a different value of β evenly sampled over a specified range. In effect, this applies different selection pressures on each LLGP population. Migrations are performed after the last adaptation step in HLGP. At this point, copies of the fittest individual in each LLGP population are sent to all the other populations, where they replace the least fit individual.

Results

Training data for five different diatomic molecules (CO, H₂, HCl, N₂, O₂) was generated using differently parameterized Morse functions, Gaussian functions and double well functions. The corresponding directed search spaces are given by:

$$F_M = D(1 - \exp(-R(r; a)))^2 + c \quad \text{Morse} \quad (10)$$

$$F_G = A \exp(R(r; a)^2) \quad \text{Gaussian} \quad (11)$$

$$F_D = D_1(1 - \exp(-R_1(r; a)))^2 + D_2(1 - \exp(-R_2(r; a)))^2 \quad \text{Double well} \quad (12)$$

Parameters D , c , A , D_1 and D_2 are optimized by including them as leaves in the trees.

The PMLGP approach was compared against a standard parallel genetic programming implementation (SPGP). In both cases, populations of 500 individuals were evolved in parallel on 8 processors for 20,000 generations. The function set $\mathcal{F} = \{+, -, \times, \div, \exp\}$ was used for internal nodes and the terminal set $\mathcal{T} = \{r, a_1, \dots, a_{10}\}$ was used for the leaf nodes.

PMLGP was shown to converge faster and achieve higher accuracy than SPGP. The obtained model of the EVB surface accurately reproduced global features of the *ab initio* data. The approach provides a basis for high quality many-body potentials for studying gas and solution phase photon reactions.

2.3 Parallel tempering

Slepoy, Peters and Thompson [47] use a hybrid approach consisting of genetic programming, Monte Carlo sampling and parallel tempering to discover the functional form of the Lennard-Jones pair potential.

Parallel tempering is an approach for parallel genetic programming where several islands (or *replicas*) evolve at a different effective temperature. High effective temperatures favor exploration by accepting new trees even if their fitness is poor, low effective temperatures favor exploitation by being sensitive to small changes in fitness. By using replicas at different temperatures the approach simultaneously performs both exploitation and exploration.

The remarkable aspect about this approach is that it marks the first large-scale application of genetic programming in materials science with interesting extensions to the canonical Koza-style algorithm and without restrictions of the hypothesis space.

The training data used consists of 10 nuclear configurations of 10 particles placed in 3-d space. The Lennard-Jones potential describes the interactions between pairs of particles, therefore a nuclear configuration's energy is given by the sum of pairwise potentials:

$$E_{\text{conf}} = \sum_{\langle i,j \rangle} V_{\text{LJ}}(r_{ij}) \quad (13)$$

where $r_{ij} = \|\mathbf{r}_i, \mathbf{r}_j\|$ is the distance between particles i and j . Fitness is defined as the negative mean squared error.

The evolutionary search is organized as a three-stage process consisting of: generation, mutation and testing.

Offspring individuals are tested for acceptance into the new population. A new tree is unconditionally accepted if its fitness exceeds the old one at the same index. Otherwise, it is accepted with the Boltzmann probability:

$$P_{\text{accept}} = \min \left\{ 1, \exp \left(\frac{F_{\text{new}} - F_{\text{old}}}{T} \right) \right\}$$

where F_{old} and F_{new} are the old and new fitness values, and T is the effective temperature.

After each generation, each sub-population exchanges one tree with its left neighbour in temperature space and one tree with its right neighbor. The trees to be swapped are selected with equal probability from their respective populations. The tree swap is accepted with a probability based on the relative Boltzmann weights of the two trees.

$$P_{\text{acc}} = \min \left\{ 1, \exp \left[\left(\frac{1}{T_i} - \frac{1}{T_{i+1}} \right) (F_{i+1} - F_i) \right] \right\}$$

Results

A large scale experiment was performed on a cluster made of 100 AMD Opteron 2.2 Ghz processors. The trees were restricted to minimum depth 3 and maximum depth 4. 200 replicas with temperatures distributed logarithmically from 0.1 to 10 were used. The replica size was chosen to be either $N = 10,000$ or $N = 50,000$ individuals. The primitive set consists of elementary operations $\mathcal{P} = \{+, -, \times, \div, \exp, |\cdot|\}$.

The proposed approach successfully discovered the Lennard-Jones potential or arithmetic equivalents within 100 generations. Interestingly, the expended effort was estimated to be somewhere in the range of 10^9 evaluated trees, which represents only a small fraction of the possible trees with depth 4 (around 2.9×10^{36}) [47].

A number of ideas for improving the physical fidelity of the developed functional forms and their generality and transferability are suggested

- Inclusion of additional properties and forces on individual atoms in the training set
- Primitive set extension to include three-body interactions
- Integration of physical knowledge (inclusion of symmetries, invariances)

2.4 Symbolically-Regressed Table KMC

In order to increase the time scale of simulations, molecular dynamics can be combined with kinetic (dynamic) Monte Carlo (KMC) techniques [9] that coarse-grain the state space, for example via discretization (e.g. assign an atom to a lattice site). The main assumption is that multiscale modeling requires only relevant information at the appropriate length or time scale.

KMC constructs a look-up table consisting of an *a priori* list of events such as atomic jumps or off-lattice jumps. This yields several order of magnitude increases in simulated time and allows to directly model many processes unapproachable by MD alone. However, identifying barrier energies from a list of events is difficult and restricts the applicability of the method.

Here, symbolic regression is proposed to identify the functional form of the potential energy surface at barrier energy points from a limited set of *ab initio* training data. The method entitled Symbolically-Regressed Table KMC (sr-KMC) [44] provides a machine learning replacement for the look-up table in KMC, thus removing the need of explicit calculation of all activation barriers.

Sastry [44] show that symbolic regression allows atomic-scale information (diffusion barriers on the potential energy surface) to be included in a long-time kinetic simulation without maintaining a detailed description of the all atomistic physics, as done within molecular dynamics.

In this approach, fitness is computed as a weighted mean absolute error between the predicted and calculated barriers, for N random configurations:

$$F = \frac{1}{N} \sum_{i=1}^N w_i |\Delta E_{\text{pred}}(\mathbf{x}_i) - \Delta E_{\text{calc}}(\mathbf{x}_i)| \quad (14)$$

Setting $w_i = |\Delta E_{\text{calc}}|^{-1}$ gives preference to predicting accurately lower-energy (most significant) events over higher energy events.

The algorithm uses the *ramped-half-and-half* tree creation method, tournament selection and Koza-style subtree crossover, subtree mutation and point mutation [34].

Results

sr-KMC is applied to the problem of vacancy-assisted migration on the surface of phase-separating $\text{Cu}_x\text{Co}_{1-x}$ at a concentrated alloy composition ($x = 0.5$). Two types of potentials (Morse and TB-SMA) are used to generate the training data via molecular dynamics. The number of active configurations is limited knowing that only atoms in the environment locally around vacancy and migrating atoms significantly influence the barrier energies.

The inline barrier function is represented from the primitive set $\mathcal{P} = \mathcal{F} \cup \mathcal{T}$, with $\mathcal{F} = \{+, -, \times, \div, \text{pow}, \text{exp}, \text{sin}\}$ and $\mathcal{T} = \{\mathbf{x}, \mathcal{R}\}$. Here, \mathbf{x} represents the current active configuration and \mathcal{R} is an ephemeral random constant.

The results show that GP predicts all barriers within 0.1% error while using less than 3% of the active configurations for training. This leads to a significant scale-up in real simulation time and a significant reduction in the CPU time needed for KMC. sr-KMC is also compared against the basic KMC approach (using a table look-up) where it was shown to perform orders of magnitude faster.

The authors note that standard basis-set regression methods are generally not competitive to GP due to the inherent difficulty in choosing appropriate basis functions and show that quadratic and cubic polynomials perform worse in terms of accuracy (within 2.5% error) while requiring energies for $\sim 6\%$ of the active configurations.

They also note that GP is robust to changes in the configuration set, the order in which configurations are used or the labeling scheme used to convert the configuration into a vector of inputs.

2.5 Hierarchical Fair Competition

Brown, Thompson and Schultz [11, 12] are able to rediscover the functional forms of known two- and three-body interatomic potentials using a parallel approach to genetic programming with extensions towards better generalization. Their implementation is based on Hierarchical Fair Competition (HFC) by Jianjun et al. [28].

The HFC framework [28] is designed towards maintaining a continuous supply of fresh genetic diversity in the population and protecting intermediate individuals who have not reached their evolutionary potential from being driven to extinction by unfair competition. It implements these goals with the help of a hierarchical population structure where individuals only compete with other individuals of similar fitness.

Brown et al. note that a correlation-based fitness measure would increase the efficiency of the search and propose the following formula using the Pearson correlation coefficient:

$$F = \frac{N}{N + 100 - 100 \left| \sum_{i=1}^N \frac{(y_i - \bar{y})(\hat{y}_i - \bar{\hat{y}})}{p_i \sigma_y \cdot p_i \sigma_{\hat{y}}} \right|} \quad (15)$$

Here, N is the number of configurations and p_i is the number of terms in the summation over g (see Equation 1). Ordinary least squares is then used to fit the prediction \hat{y} to the data by introducing scale and intercept terms to the functions g and h :

$$E = \sum_{\langle i,j \rangle} (a \cdot g(\mathbf{r}_i, \mathbf{r}_j) + b) + \sum_{\langle i,j,k \rangle} (c \cdot h(\mathbf{r}_i, \mathbf{r}_j, \mathbf{r}_k) + d) \quad (16)$$

The approach is implemented in `PM-DREAMER`, an open-source software package developed on top of the Open Beagle library for evolutionary computation [19], using its available genetic operators. These include several mutations (standard, shrink, swap, constant), subtree-swapping crossover, tournament selection and elitism:

- *Standard* mutation replaces a node in the tree with a randomly-generated subtree.
- *Swap* mutation swaps two nodes in the tree.
- *Shrink* mutation replaces a subtree with one of its arguments.
- *Swap subtree* mutation swaps a subtree's arguments.
- *Ephemeral* mutation changes the value of a constant in the tree.

Additionally, `PM-DREAMER` implements support for distributed evolution using the MPI standard and introduces migration operators that exchange individuals between sub-populations at fixed intervals.

Bloat reduction strategies are implemented to prevent the expression trees from becoming increasingly large, a tendency observed especially in the case of three-body modeling. Two strategies are tested:

- Using a simplification operator which replaces subtrees that evaluate to a constant value with the constant value. This operator is applied generationally at a fixed interval.
- Using penalty terms to the fitness function: in this case the fitness is decreased based on a threshold penalty size value s_b and a maximum penalty size s_e , such that trees with length $< s_b$ are not penalized at all, and trees with length $> s_e$ are penalized fully (fitness is set to zero).

Local search. Local search based on the derivative-free Nelder-Mead simplex algorithm is employed with a set probability, optimizing either a single constant or all the constants in the expression.

HFC Extension. Brown et al. implement HFC in a parallel manner by allowing populations with different fitness thresholds to evolve in parallel, with periodic migrations between them. After migrations, populations that grow too large are “decimated” by removal of the least fit individuals, while populations that grow too small are supplemented with new randomly-generated individuals.

The population fitness thresholds are adapted during the search using two strategies: the first strategy uses a percentile parameter p which determines the fitness threshold such that p percent of individuals have equal or lower fitness. The second strategy uses fixed thresholds determined by the first non-zero threshold along with a scaling parameter equal to the ratio between successive thresholds.

Results

Training data for two- and three-body interactions was generated using the Lennard-Jones and Stillinger-Weber potentials (see Section 5 Eqs. 18, 20). In both cases, five configurations were used for training and 50 configurations were used for testing, in order to realistically represent the problem of obtaining models for condensed phases from a small training set. The generated data includes pairwise distances between atoms, the energy and the force on a single atom.

The authors compare a standard parallel island-based evolutionary model against a parallel HFC evolutionary model, using 32 islands with a population of 10 000 individuals each, evolved over a period of 100 generations.

The primitive set used was $\mathcal{P} = \{+, -, \times, \div, \text{pow}, \text{exp}, \text{log}, |\cdot|\}$, tournament selection was used with a tournament size of 6 and, in the case of the standard evolutionary model, 500 individuals were migrated between islands every 5 generations. For the HFC evolutionary model, the migration took place every generation, the first fitness threshold was set to 0.1 and the threshold ratio was set to 1.0. A detailed description of the other algorithm parameters is given in [12].

After an initial tuning phase, the authors note that the number of interactions per energy point greatly increases the runtime requirements for optimization. The C++ implementation of PM-DREAMER is capable of doing vectorized evaluation of two- and three-body interatomic potential models using SIMD instructions in a manner similar to batched tree interpretation more typically used in GP. With vectorization, the evolutionary algorithm was found to perform roughly four times faster.

Local search was performed with varying probability on all constants in an expression using a maximum of 6 iterations of the Nelder-Mead algorithm. Simplification is performed every 20 generations.

Overall, the authors show that the HFC strategy consistently outperforms the standard generational evolutionary strategy and is able to find very accurate approximations for the targeted empirical potentials (Lennard-Jones and Stillinger-Weber).

2.6 Potential Optimization by Evolutionary Techniques (POET)

POET [24] distinguishes itself from previously described approaches through an extended primitive set which includes summation symbols that aggregate local energy values around each atom, smoothing functions meant to exploit the “short-sightedness” of atomic interactions as well as leaf nodes representing the atomic neighborhood interaction radius.

The primitive set used by the algorithm consists of the function set $\mathcal{F} = \{\sum, f, +, -, \times, \div, \text{pow}\}$ and terminal set $\mathcal{T} = \{\mathcal{R}, r\}$. Here, \sum are summation symbols, f are smoothing symbols, \mathcal{R} represents an ephemeral constant and, like before, r represents the distance between atoms. Distances are considered within the neighborhood of each atom according to inner and outer cutoff radii r_{in} and r_{out} .

An exemplary POET-tree including the special symbols \sum , f and r is shown in Figure 2. This tree corresponds to the following function which returns the predicted value of the local energy E_i around the i th atom considering the distances r_{ij} to its neighbors:

$$E_i = 7.51 \sum_j r_{ij}^{3.98-3.93r_{ij}} f(r_{ij}) + \left(28.01 - 0.03 \sum_j r_{ij}^{11.73-2.93r_{ij}} f(r_{ij}) \right) \left(\sum_j f(r_{ij}) \right)^{-1}$$

Hernandez et al. employ a parallel version of genetic programming where twelve populations are evolved simultaneously. The recombination pool in each population is filled from three separate sets of models: a set from the current population, a global set maintained with the overall best (non-dominated) individuals with regard to fitness, complexity and execution speed, and a set of individuals from the other populations. These sets are periodically filled up with individuals at preset intervals.

New individuals are generated by means of crossover and mutation. Crossover replaces a random subtree in the root parent with either another random subtree from another parent or with a linear combination of random subtrees from two different parents. The first method was applied with probability 0.9 while the second method was applied with probability 0.1.

Mutation can replace a subtree with a randomly generated one, swap the arguments of non-commutative symbols or change the symbols of function nodes. Tree initialization is done using Koza’s ramped-half-and-half method where the tree depth is sampled from a Gaussian distribution with $\mu = 5$ and $\sigma = 1$.

Local optimization of model coefficients is performed online during the run with the help of a covariance matrix adaptation evolution strategy (CMA-ES) optimizer and a conjugate gradient (CG) optimizer. CMA-ES is used to optimize the coefficients of models in the global set every 10,000 crossover and mutation operations. The CG algorithm is used to perform one optimization step for every newly generated individual.




Fig. 2: Interatomic potential representing local energy E_i around the i th atom in electron volts. The expression resembles that of an embedded atom model with a pairwise repulsive term and a many-body attractive term formed by a non-linear transformation of pairwise interactions. Caption and image from Hernandez et al. [24].

Results

The proposed approach is validated using training data from DFT molecular dynamics simulations containing snapshots of atomic positions, energies, forces and stresses for an atomic system of 32 Cu atoms. The fitness measure is an aggregation of the energy, force and stress errors:

$$F = 1000 \cdot (0.5\text{MSE}_{\text{energy}} + 0.4\text{MSE}_{\text{force}} + 0.1\text{MSE}_{\text{stress}})$$

The authors demonstrate POET's ability to rediscover Lennard-Jones and Sutton-Chen potentials. The generated models displayed low overfitting and high generalization being able to maintain high predictive accuracy for properties on which they were not trained. The simplicity of the models allows them to predict energies with speeds in the order of microseconds per atom, about 1-4 orders of magnitude faster than other ML potentials. The authors also note that such simple models bring the additional advantage of requiring relatively small amounts of training data.

In terms of runtime performance of POET itself, the authors report 330 CPU-hours spent finding the exact Lennard-Jones potential, 3600 CPU-hours finding the exact Sutton-Chen potential and 360 CPU-hours to find the three best performing models reported in [24]. POET code is open-source and available online¹.

2.7 Other applications

Makarov and Metiu [38] use directed genetic programming to find analytic solutions to the time-independent Schrödinger equation. The training data is generated by inverting the Schrödinger equation such that the potential is a functional depending on the wave function and the energy.

Kenoufi and Kholmurodov [30] used symbolic regression to rediscover the Lennard-Jones potential and discovered a new potential for an argon dimer, using *ab initio* data from DFT simulations.

Mueller et al. [40] used symbolic regression for discovering relevant descriptors in hydrogenated nanocrystalline silicon with very low crystalline volume fraction, with applications in improving optical absorption efficiency in thin-film photovoltaics.

Wang et al. [53] used symbolic regression to discover the Johnson-Mehl-Avrami-Kolmogorov transformation kinetics law in the recrystallization process of copper; and the Landau free energy functional form for the displacive tilt transition in perovskite LaNiO_3 .

Eldridge et.al [18] used the NSGA-III algorithm to learn interatomic potentials for carbon. The approach considered training error, individual age and individual complexity as objectives and was able to find simple and accurate potential functions.

2.8 Summary discussion

State-of-the-art SR approaches for modeling interatomic potentials recognize the need for domain-specific extensions and hybridizations towards promoting physical plausibility and achieving high accuracy, while making the most out of the usually scarce quantities of available *ab initio* training data.

Several extensions and hybridizations are used to augment the classic (Koza-style) genetic programming algorithm and increase its search performance. Parallel, island-based approaches are employed in all of the discussed methods, on the one hand, to more efficiently search the hypothesis space and on the other hand, to achieve higher throughput and alleviate the high computational costs of summations over two- and three-body atomic interactions.

Physical plausibility is promoted by restricting the hypothesis space (directed search), including domain specific information into the fitness function (e.g. weighing down high-energy points) or including additional targets (forces and stresses).

¹ <https://gitlab.com/muellergroup/poet>

The results achieved so far have demonstrated the ability of symbolic regression to discover highly accurate and physically-plausible functional forms which can increase, due to their simplicity and efficiency, the performance of particle simulations (allowing them to run at larger scales or for longer times). At the same time, since the models are inherently more simple than similar black-box models such as ANNs, they tend to require a lesser amount of training data, which increases their applicability.

Overall, it can be concluded that symbolic regression represents a very promising approach for discovering more accurate and efficient potentials. However, designing evolutionary systems for this application area requires consideration of specific challenges as described in Section 1.2. In the following section we discuss several ideas towards a GP system design which is able to address the domain-specific requirements of interatomic potential.

3 Designing GP for modeling interatomic potentials

The problem of modeling interatomic potentials from data has the main particularity that data comes in the form of atomic configuration snapshots. Each configuration describes the positions of the atoms, its energy and optionally other properties (forces, stresses). Canonically, these data snapshots are generated by molecular dynamics simulation packages such as LAMMPS [42] or VASP [35] and come in specific formats, see e.g. POSCAR².

At the minimum, raw data has the following form:

$$\begin{bmatrix} \mathbf{r}_1^{(1)} & \dots & \mathbf{r}_M^{(1)} & E^{(1)} \\ \vdots & & \vdots & \vdots \\ \mathbf{r}_1^{(N)} & \dots & \mathbf{r}_M^{(N)} & E^{(N)} \end{bmatrix}$$

where $\mathbf{r}_i^{(k)}$ is the position of the i th atom in the snapshot k and $E^{(k)}$ is the associated potential energy value.

Since atomic interactions are computed based on the distances between atoms, the Cartesian coordinates need to be converted into sets of pairwise distances relative to each atom. It is then the role of the genetic system to evolve an accurate functional relationship between distances r_{ij} and potential energy. For each training sample k , the symbolic regression model needs to process a set of pairwise atomic distances into a prediction for the energy with the help of summation symbol \sum .

As it becomes apparent from studying previous approaches described in Section 2, modeling interatomic potentials is a non-trivial problem which requires substantial computational resources. Previous implementations employed different strategies for parallelism as well as other optimization techniques such as vectorized model evaluation in order to speed up the search. Additionally, most approaches employed local search in order to improve model coefficients during evolution.

² <https://www.vasp.at/wiki/index.php/POSCAR>

For this reason, we opt to extend the framework Operon [13] with additional functionality for modeling interatomic potentials. Operon already benefits from a fine-grained parallelism model designed for scalability and was shown to perform well on a variety of symbolic regression problems [14]. Additionally, it features support for local optimization using the Levenberg-Marquardt optimization algorithm, where the gradient is obtained via automatic differentiation.

We adopt a multi-objective approach based on the NSGA2 algorithm [16] where model length is used alongside prediction accuracy in order to promote parsimony, interpretability and generalization.

3.1 Symbolic regression in Operon

Operon is a C++ framework for symbolic regression that employs logical parallelism during evolution, such that every new offspring individual is generated in its own logical thread. An example evolutionary algorithm implemented in Operon as an operator graph is shown in Figure 3.

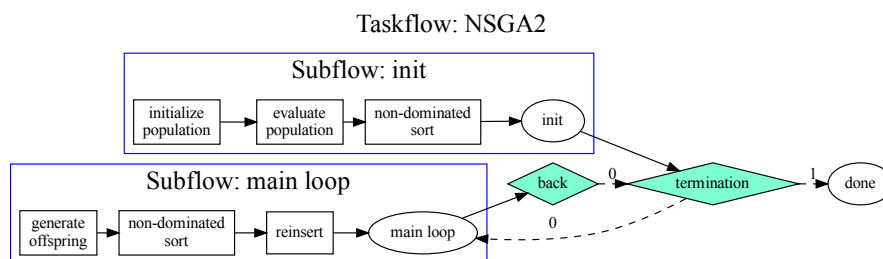


Fig. 3: Taskflow describing the NSGA2 algorithm in Operon. Each individual task within the subflows (initialization, evaluation, offspring generation) executes in parallel, using a number of logical threads equal to the population size.

Operon uses a linear encoding where each tree is represented as a postfix sequence of nodes. Each node has typical attributes such as length, depth, arity or opcode. Evaluation efficiency is achieved by employing a batched tree interpreter, which iterates over the tree nodes and executes the corresponding functions on fixed-size batches of data. As the batch size is known at compile-time, these operations are vectorized. The entire tree evaluation infrastructure relies on the *Eigen C++* library [22] for efficient, vectorized execution.

3.1.1 Implementing the Σ symbol

The tree interpreter represents a generic approach to tree evaluation and is agnostic of the actual primitive set used by the algorithm. Each node is mapped to a callable³ (stateful function object) which defines the functional transformation. The callables themselves are required to satisfy a certain function signature and to operate in both scalar and dual number domains.

This mechanism facilitates the extension of the default primitive set with any kind of ad hoc functionality – the Σ (summation) symbol in this particular application. Figure 4 shows the general workflow for processing a set of atomic positions into pairwise distances and using them to estimate the potential energy. The function F represents a symbolic functional form which includes Σ symbols over pairwise atomic distances. Since the Σ symbol is essentially a reduction operator⁴, the actual number of dimensions of the input data is three: N snapshots, M atoms, L pairwise distances (where L dynamically depends on the cut-off radius r_{out}).

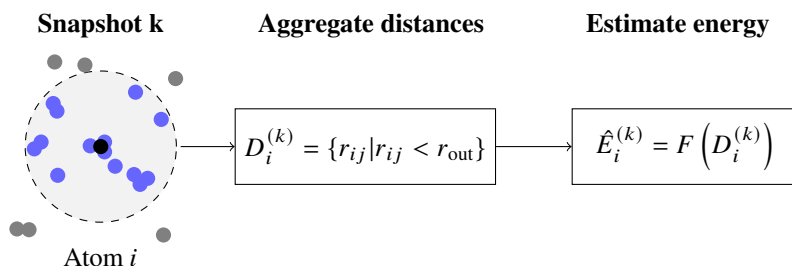


Fig. 4: Prediction of atom energies using SR. The total energy is then $\hat{E}^{(k)} = \sum_i \hat{E}_i^{(k)}$.

Like many other evolutionary frameworks, Operon relies on a dataset object which holds tabular data in two dimensions: X features \times Y observations. Therefore, it is not straightforward to accommodate an additional data dimension without significant redesign work. However, it is relatively easy to incorporate an extra, inner dataset into the function object associated to the Σ symbol, which will contain the values in the third dimension (interatomic distances).

For this mechanism to work, a convention is necessary: the *outer* dataset will contain the target energy values as well as an input variable r whose value is always 1 (this value was chosen arbitrarily as a non-problematic constant which does not cause discontinuities). The variable r simply acts a placeholder for the pairwise atomic distance values. The *inner* dataset will contain the actual pairwise distances under the same input name r . The distances are computed when the atomic coordinate values are loaded into the callable. A nested tree interpreter is then used to evaluate the

³ https://en.wikipedia.org/wiki/Callable_object

⁴ https://en.wikipedia.org/wiki/Reduction_operator

current Σ -subtree using the inner dataset as input. Similar to Hernandez et al. [24], the Σ symbol also applies a smoothing function on its output (see Equation (7) in [24]) with the inner and outer cutoff radii equal to 3\AA and 5\AA , respectively.

Under this set of rules, a leaf node corresponding to the input variable r will evaluate to 1 when not under a Σ symbol, and to the set of pairwise distance values corresponding to the current atomic configuration otherwise. Additionally, to disallow nesting of Σ symbols, the behavior is dynamically switched depending on the surrounding tree context: if a Σ symbol finds itself under another Σ symbol, then it simply acts as the identity function $f(x) = x$. This convention does not impact the evolutionary system’s ability to discover interatomic potential functional forms.

3.2 Empirical validation

We demonstrate the capabilities of the proposed NSGA2-based multi-objective approach using the *ab initio* data used by Hernandez et al. [24]. This data consists of 150 snapshots of 32-atom DFT molecular dynamics simulations of copper (Cu): 50 snapshots at 300 K (NVT), 50 snapshots at 1400 K (NVT) and 50 snapshots at 1400 K (NPT at 100 kPa). Although the data also contains components of forces and virial stress tensors, only the energy was used as a modeling target in this experiment. The data consisting of 150 configurations is shuffled and split equally into training and test partitions.

3.2.1 Experimental setup

The experiment used a fixed set of parameters shown in Table 1. The primitive set was varied and consisted of different symbol combinations, as shown in Table 2: with and without the power function, and alternating between \div and aq , where $\text{aq}(x) = \sqrt{x^2 + 1}$.

Two input variables are used: r as a placeholder for atomic distances and $q = \frac{1}{r}$ as a placeholder for the inverse of r , given that some empirical potentials like Lennard-Jones explicitly use the inverse in their formula. Each experimental configuration was repeated 50 times and the median values were reported (with the exception of runtime, which was averaged). Errors are reported as median \pm standard deviation. Model length was computed as the length of the simplified representation returned by *Sympy*, using the infix textual representation of the best individual as input.

3.2.2 Results

Results aggregated over 50 runs for each configuration are shown in Table 2, alongside p-value matrices computed using the Kruskal test. Significance is encoded in Tables 3 and 4 using font weight and color: values lower than $\alpha = 0.01$ are shown in bold

population size	10,000 individuals
tree initialization	balanced tree creator (BTC) [13]
max tree length	20
max tree depth	10
crossover probability	100%
crossover operator	subtree crossover
mutation probability	25%
mutation operator	uniformly chosen from: <ul style="list-style-type: none"> • subtree removal/insertion/replacement • change function symbol • change variable name • additive one point leaf mutation ($v = v + \mathcal{N}(0, 1)$) • discrete point leaf mutation ($v \leftarrow \text{math constant: } \pi, e, \dots$)
selection operator	crowded tournament selection, group size = 17
objectives	Pearson R^2 and model length
evaluation budget	10^8 fitness evaluations

Table 1: NSGA2 parameters

black font, values lower than $\alpha = 0.05$ are shown in black font, while all the other values are shown in gray. The direction of the relationship is determined using a comparison of median values and shown as \uparrow (worse/higher error) or \downarrow (better/lower error) symbols prefixed to the values.

The overall best models from all runs and all configurations are shown in Table 5. These models have been selected based on both test accuracy and simplicity of their functional form. Two other models with better test score have been discarded due to very complex structure or very large coefficient values. Table 5 illustrates this fact by displaying the absolute rank of each model (based purely on test accuracy and disregarding other criteria).

Interestingly, the arithmetic-only configurations A, B, C generated 4 out of 5 of the selected best models. Although configuration A produced significantly worse ($p < 0.01$) training accuracies than all other configurations, it did not produce worse models in terms of generalization, where it is only worse than E. Nevertheless, the explicit inclusion of $1/r$ as an input seems to help the search.

It is also worth noting that configurations using the analytic quotient instead of (unprotected) division generally perform better on the training data ($p < 0.05$), but do not perform better on the test data. For example, configuration K is better than A, B, C, H, I in terms of training accuracy, but is not better than any of them in terms of test (on the contrary, it is worse than E at $p < 0.05$). From this we can surmise that in this particular test setting and for this particular data, AQ does not offer an advantage compared to normal division.

Overall, judging from median error values and statistical significance p-values, there is no clear winner among the tested configurations. However, a pattern emerges when observing the functional forms of the best models, mostly originating from configurations B and C. After simplification using *Sympy*, the models become highly similar with the same mathematical structure consisting of a sum of three factors in the numerator (each including the inverse of r) and another sum in the denominator

(also including the inverse of r). Although these models are remarkably simple, further testing is required to validate their properties and behavior.

In terms of runtime, the proposed approach is efficient, with the longest run taking on average 290 seconds to evolve a population of 10,000 individuals for 1000 generations on a single multicore computer. In comparison, Hernandez et al. [24] report 360 CPU-hours expended on finding accurate GP models.

Table 2: Operon NSGA2 Results

Table 3: Training error p-value matrix using the Kruskal statistical test. Significance shown by bold black font ($p < 0.01$), black font ($p < 0.05$) or gray (no significance). Relationship direction given by comparison of medians: \uparrow (worse/higher error), \downarrow (better/lower error).

Table 4: Test error p-value matrix using the Kruskal statistical test. Significance shown by bold black font ($p < 0.01$), black font ($p < 0.05$) or gray (no significance). Relationship direction given by comparison of medians: \uparrow (worse/higher error), \downarrow (better/lower error).

Table 5: Overall best models, where ID identifies the configuration in Table 2.

4 Conclusion

This work surveyed the main applications of SR in Materials Science, namely for the discovery of simple and efficient models of interatomic potentials. Both previous results, as well as results obtained by our own proposed approach and described in this paper, suggest that SR is capable of finding accurate models that can further the capabilities of particle simulations.

Similar to POET [24], our approach does not restrict the search space in any way (with the exception of tree length and depth limits) and is therefore capable of finding models that do not resemble previously known, empirical potential functions. At the same time, should a directed search be required, the framework is trivial to extend with this feature.

Empirical testing shows that relatively simple primitive sets are powerful enough to discover accurate potential functions with good extrapolation behavior. On this data, no advantage was found in using the analytical quotient over standard division. More experiments will be required to establish the benefits of larger primitive sets, for example ones that include logarithmic, exponential or trigonometric functions.

Several other aspects like a more comprehensive search in the space of hyper-parameters or an exploration of the effects of local search also need to be fully investigated in the future. Compared to other works described in our survey, our approach did not diverge from the “vanilla” version of GP, using a classical multi-objective approach (NSGA2) together with a domain specific primitive set. It will be also worthwhile to explore various ways to scale up the search using multiple populations and more sophisticated evolutionary models.

Future development directions include expanding the capabilities of the framework to include three- or many-body interactions, to consider model derivatives in order model atomic forces as well, and overall to improve its ability to incorporate and respect the fundamental laws of this kind of physical systems.

5 Appendix

Empirical potentials

For a comprehensive overview of empirical potentials we recommend the work of Araújo and Ballester [2]. Below we give a casual overview of the most important empirical potentials mentioned in this contribution.

Morse potential

This is an empirical potential used to model diatomic molecules.

$$V_M(r) = D \left(1 - \exp(-a(r - r_0))^2 \right) \quad (17)$$

where D is the dissociation energy, r is the distance between atoms, a is a set of parameters and r_0 is the equilibrium bond distance.

Lennard-Jones potential

The Lennard-Jones potential models soft repulsive and attractive interactions and can describe electronically neutral atoms or molecules. Interacting particles repel each other at very close distance, attract each other at moderate distance, and do not interact at infinite distance.

$$V_{LJ}(r) = 4\varepsilon \left[\left(\frac{\sigma}{r} \right)^{12} - \left(\frac{\sigma}{r} \right)^6 \right] \quad (18)$$

where r is the distance between atoms, ε is the dispersion energy and σ is the distance at which the particle-particle potential energy V is zero.

Lippincott potential

Lippincott [48] potential involves an exponential of interatomic distances

$$V_{LIP}(r) = D \left(1 - \exp\left(\frac{-n(r - r_0)^2}{2r}\right) \right) (1 + aF(r)) \quad (19)$$

where D is the dissociation energy, r is the distance between atoms, r_0 is the equilibrium bond distance and a and n are parameters. $F(r)$ is a function of internuclear distance such that $F(r) = 0$ when $r = \infty$ and $F(r) = \infty$ when $r = 0$.

Stillinger-Weber potential

The Stillinger-Weber potential [49] models two- and three-body interactions by taking into account not only the distances between atoms but also the bond angles:

$$V_{\text{SW}}(r) = \sum_{\langle i,j \rangle} \phi_2(r_{ij}) + \sum_{\langle i,j,k \rangle} \phi_3(r_{ij}, r_{ik}, \theta_{ijk}) \quad (20)$$

where

$$\phi_2(r_{ij}) = A\varepsilon \left[B \left(\frac{\sigma}{r_{ij}} \right)^p - \left(\frac{\sigma}{r_{ij}} \right)^q \right] \exp \left(\frac{\sigma}{r_{ij} - a\sigma} \right) \text{ and} \quad (21)$$

$$\phi_3(r_{ij}, r_{ik}, \theta_{ijk}) = \lambda\varepsilon [\cos \theta_{ijk} - \cos \theta_0]^2 \times \exp \left(\frac{\gamma\sigma}{r_{ij} - a\sigma} \right) \exp \left(\frac{\gamma\sigma}{r_{ik} - a\sigma} \right) \quad (22)$$

Sutton-Chen potential

The Sutton-Chen potential [50] has been used in molecular dynamics and Monte Carlo simulations of metallic systems. It offers a reasonable description of various bulk properties, with an approximate many-body representation of the delocalized metallic bonding:

$$V_{\text{SC}} = \sum_{\langle i,j \rangle} U(r_{ij}) - \sum_i u\sqrt{\rho_i} \quad (23)$$

Here, the first term represents the repulsion between atomic cores and the second term models the bonding energy due to the electrons. Both terms are further defined in terms of reciprocal power so that the complete expression is:

$$V_{\text{SC}} = \epsilon \left[\sum_{\langle i,j \rangle} \left(\frac{a}{r_{ij}} \right)^n - C \sum_i \sqrt{\sum_j \left(\frac{a}{r_{ij}} \right)^m} \right] \quad (24)$$

where C is a dimensionless parameter, ϵ is a parameter with dimensions of energy, a is the lattice constant, m, n are positive integers with $n > m$ and r_{ij} is the distance between the i th and j th atoms.

References

- [1] A. Agrawal and A. Choudhary. Perspective: Materials informatics and big data: Realization of the “fourth paradigm” of science in materials science. *APL Materials*, 4(5):053208, 2016.
- [2] Judith P. Araújo and Maikel Y. Ballester. A comparative review of 50 analytical representation of potential energy interaction for diatomic systems: 100 years of history. *International Journal of Quantum Chemistry*, 121(24):e26808, 2021.

- [3] James E. Baker. Reducing bias and inefficiency in the selection algorithm. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic Algorithms and Their Application*, pages 14–21, USA, 1987. L. Erlbaum Associates Inc.
- [4] R. M. Balabin and E. I. Lomakina. Support vector machine regression (ls-svm)—an alternative to artificial neural networks (anns) for the analysis of quantum chemistry data? *Phys. Chem. Chem. Phys.*, 13:11710–11718, 2011.
- [5] A. P. Bartók, R. Kondor, and G. Csányi. On representing chemical environments. *Phys. Rev. B*, 87:184115, May 2013.
- [6] J. Behler. Perspective: Machine learning potentials for atomistic simulations. *J. Chem. Phys.*, 145(17):170901, 2016.
- [7] Michael A. Bellucci and David F. Coker. Empirical valence bond models for reactive potential energy surfaces: A parallel multilevel genetic program approach. *The Journal of Chemical Physics*, 135(4):044115, 2011.
- [8] Michael A. Bellucci and David F. Coker. Molecular dynamics of excited state intramolecular proton transfer: 3-hydroxyflavone in solution. *The Journal of Chemical Physics*, 136(19):194505, 2012.
- [9] K. Binder, D. Heermann, Lyle Roelofs, A. John Mallinckrodt, and Susan McKay. Monte carlo simulation in statistical physics. *Computers in Physics*, 7(2):156–157, 1993.
- [10] A. Brown, A. B. McCoy, B. J. Braams, Z. Jin, and J. M. Bowman. Quantum and classical studies of vibrational motion of CH_5^+ on a global potential energy surface obtained from a novel ab initio direct dynamics approach. *J. Chem. Phys.*, 121(9):4105–4116, 2004.
- [11] Michael W. Brown, Aidan P. Thompson, Jean-Paul Watson, and Peter A. Schultz. Bridging scales from ab initio models to predictive empirical models for complex materials. Technical report, Laboratories, Sandia National, 2008.
- [12] W. M. Brown, A. P. Thompson, and P. A. Schultz. Efficient hybrid evolutionary optimization of interatomic potential models. *J. Chem. Phys.*, 132(2):024108, 2010.
- [13] Bogdan Burlacu, Gabriel Kronberger, and Michael Kommenda. Operon C++: An efficient genetic programming framework for symbolic regression. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference Companion*, GECCO '20, pages 1562–1570, internet, July 8-12 2020. Association for Computing Machinery.
- [14] William G. La Cava, Patryk Orzechowski, Bogdan Burlacu, Fabrício Olivetti de França, Marco Virgolin, Ying Jin, Michael Kommenda, and Jason H. Moore. Contemporary symbolic regression methods and their relative performance. *CoRR*, abs/2107.14351, 2021.
- [15] R. Chen, K. Shao, B. Fu, and D. H. Zhang. Fitting potential energy surfaces with fundamental invariant neural network. ii. generating fundamental invariants for molecular systems with up to ten atoms. *J. Chem. Phys.*, 152(20):204307, 2020.
- [16] Kalyanmoy Deb, Samir Agrawal, Amrit Pratap, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE Trans. Evol. Comput.*, 6(2):182–197, 2002.

- [17] P. O. Dral. Quantum chemistry in the age of machine learning. *J. Phys. Chem. Lett.*, 11(6):2336–2347, 2020. PMID: 32125858.
- [18] Andrew Eldridge, Alejandro Rodriguez, Ming Hu, and Jianjun Hu. Genetic programming-based learning of carbon interatomic potential for materials discovery, 2022.
- [19] Christian Gagné and Marc Parizeau. Genericity in evolutionary computation software tools: Principles and case study. *International Journal on Artificial Intelligence Tools*, 15(2):173–194, April 2006.
- [20] H. Gao, J. Wang, and J. Sun. Improve the performance of machine-learning potentials by optimizing descriptors. *J. Chem. Phys.*, 150(24):244110, 2019.
- [21] L. M. Ghiringhelli, J. Vybiral, S. V. Levchenko, C. Draxl, and M. Scheffler. Big data of materials science: Critical role of the descriptor. *Phys. Rev. Lett.*, 114:105503, Mar 2015.
- [22] Gaël Guennebaud, Benoît Jacob, et al. Eigen v3. <http://eigen.tuxfamily.org>, 2010.
- [23] C. M. Handley and J. Behler. Next generation interatomic potentials for condensed systems. *European Physical Journal B*, 87(7):152, July 2014.
- [24] A. Hernandez, A. Balasubramanian, F. Yuan, S. A. M. Mason, and T. Mueller. Fast, accurate, and transferable many-body interatomic potentials by symbolic regression. *NPJ Computational Materials*, 5(1):112, 2019.
- [25] T. Hey, K. Butler, S. Jackson, and J. Thiyagalingam. Machine learning and big scientific data. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 378(2166):20190054, 2020.
- [26] Lauri Himanen, Amber Geurts, Adam Stuart Foster, and Patrick Rinke. Data-driven materials science: Status, challenges, and perspectives. *Advanced Science*, 6(21):1900808, 2019.
- [27] Adam Hospital, Josep Ramon Goñi, Modesto Orozco, and Josep L Gelpí. Molecular dynamics simulations: advances and applications. *Advances and applications in bioinformatics and chemistry: AABC*, 8:37, 2015.
- [28] Jianjun Hu, Erik Goodman, Kisung Seo, Zhun Fan, and Rondal Rosenberg. The hierarchical fair competition (hfc) framework for sustainable evolutionary algorithms. *Evolutionary Computation*, 13(2):241–277, 06 2005.
- [29] J. Ischtwan and M. A. Collins. Molecular potential energy surfaces by interpolation. *J. Chem. Phys.*, 100(11):8080–8088, 1994.
- [30] Abdelouahab Kenoufi and Kholmurzo Kholmurodov. Symbolic regression of interatomic potentials via genetic programming. *Biol. Chem. Res*, 2:1–10, 2015.
- [31] Chiho Kim, Ghanshyam Paliana, and Ramamurthy Ramprasad. From organized high-throughput data to phenomenological theory using machine learning: The example of dielectric breakdown. *Chemistry of Materials*, 28(5):1304–1311, 2016.
- [32] K. H. Kim, Y. S. Lee, T. Ishida, and G-H Jeung. Dynamics calculations for the lih+h li+h2 reactions using interpolations of accurate ab initio potential energy surfaces. *J. Chem. Phys.*, 119(9):4689–4693, 2003.
- [33] W. Kohn and L. J. Sham. Self-consistent equations including exchange and correlation effects. *Phys. Rev.*, 140:A1133–A1138, Nov 1965.

- [34] John R. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge, MA, USA, 1992.
- [35] G. Kresse and J. Furthmüller. Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set. *Phys. Rev. B*, 54:11169–11186, Oct 1996.
- [36] Aaron Kusne, Tim Mueller, and Ramamurthy Ramprasad. Machine learning in materials science: Recent progress and emerging applications. *Reviews in Computational Chemistry*, 2016-05-06 2016.
- [37] D. E. Makarov and H. Metiu. Fitting potential-energy surfaces: A search in the function space by directed genetic programming. *J. Chem. Phys.*, 108(2):590–598, 1998.
- [38] Dmitrii E. Makarov and Horia Metiu. Using genetic programming to solve the schrödinger equation. *The Journal of Physical Chemistry A*, 104(37):8540–8545, 2000.
- [39] T. Mueller, A. Hernandez, and C. Wang. Machine learning for interatomic potential models. *J. Chem. Phys.*, 152(5):050902, 2020.
- [40] T. Mueller, E. Johlin, and J. C. Grossman. Origins of hole traps in hydrogenated nanocrystalline and amorphous silicon revealed through machine learning. *Phys. Rev. B*, 89:115202, 2014.
- [41] Ghanshyam Pilania. Machine learning in materials science: From explainable predictions to autonomous design. *Computational Materials Science*, 193:110360, 2021.
- [42] Steve Plimpton. Fast parallel algorithms for short-range molecular dynamics. *Journal of Computational Physics*, 117(1):1–19, 1995.
- [43] T. Rothe, J. Schuster, F. Teichert, and E.E. Lorenz. *Machine Learning Potentials - State of the Research and Potential Applications for Carbon Nanostructures*. Technische Universität, Faculty of Natural Sciences, Institute of Physics, 2019.
- [44] Kumara Narasimha Sastry. *Genetic algorithms and genetic programming for multiscale modeling: Applications in materials science and chemistry and advances in scalability*. PhD thesis, University of Illinois, Urbana-Champaign, March 2007.
- [45] K. Shao, J. Chen, Z. Zhao, and D. H. Zhang. Communication: Fitting potential energy surfaces with fundamental invariant neural network. *J. Chem. Phys.*, 145(7):071101, 2016.
- [46] A. V. Shapeev. Moment tensor potentials: A class of systematically improvable interatomic potentials. *Multiscale Modeling & Simulation*, 14(3):1153–1173, 2016.
- [47] A. Slepoy, M. D. Peters, and A. P. Thompson. Searching for globally optimal functional forms for interatomic potentials using genetic programming with parallel tempering. *J. Comput. Chem.*, 28(15):2465–2471, 2007.
- [48] Derek Steele, Ellis R. Lippincott, and Joseph T. Vanderslice. Comparative study of empirical internuclear potential functions. *Rev. Mod. Phys.*, 34:239–251, Apr 1962.
- [49] Frank H. Stillinger and Thomas A. Weber. Computer simulation of local order in condensed phases of silicon. *Phys. Rev. B*, 31:5262–5271, Apr 1985.

- [50] A. P. Sutton and J. Chen. Long-range finnis–sinclair potentials. *Philosophical Magazine Letters*, 61(3):139–146, 1990.
- [51] A.P. Thompson, L.P. Swiler, C.R. Trott, S.M. Foiles, and G.J. Tucker. Spectral neighbor analysis method for automated generation of quantum-accurate interatomic potentials. *J. Comput. Phys.*, 285:316–330, 2015.
- [52] Oliver T. Unke, Stefan Chmiela, Huziel E. Saucedo, Michael Gastegger, Igor Poltavsky, Kristof T. Schütt, Alexandre Tkatchenko, and Klaus-Robert Müller. Machine learning force fields. *Chemical Reviews*, 0(0):null, 2021. PMID: 33705118.
- [53] Yiqun Wang, Nicholas Wagner, and James M Rondinelli. Symbolic regression in materials science. *MRS Communications*, 9(3):793–805, 2019.
- [54] L. Zhang, J. Han, H. Wang, R. Car, and E. Weinan. Deep potential molecular dynamics: A scalable model with the accuracy of quantum mechanics. *Phys. Rev. Lett.*, page 143001, 2018.