

PAPER • OPEN ACCESS

## Massively parallel fitting of Gaussian approximation potentials

To cite this article: Sascha Klawohn *et al* 2023 *Mach. Learn.: Sci. Technol.* **4** 015020

View the [article online](#) for updates and enhancements.

### You may also like

- [Computational study of crystal defect formation in Mo by a machine learning molecular dynamics potential](#)  
F J Domínguez-Gutiérrez, J Byggmästar, K Nordlund et al.
- [Variational Gaussian approximation for Poisson data](#)  
Simon R Arridge, Kazufumi Ito, Bangti Jin et al.
- [Asymptotic approximations for the phase space Schrödinger equation](#)  
Panos D Karageorge and George N Makrakis



## PAPER

## Massively parallel fitting of Gaussian approximation potentials

## OPEN ACCESS

RECEIVED  
19 October 2022REVISED  
15 November 2022ACCEPTED FOR PUBLICATION  
29 November 2022PUBLISHED  
16 February 2023Sascha Klawohn<sup>1</sup> , James R Kermode<sup>1</sup> and Albert P Bartók<sup>1,2,\*</sup> <sup>1</sup> Warwick Centre for Predictive Modelling, School of Engineering, University of Warwick, Coventry CV4 7AL, United Kingdom<sup>2</sup> Department of Physics, University of Warwick, Coventry CV4 7AL, United Kingdom

\* Author to whom any correspondence should be addressed.

E-mail: [apbartok@gmail.com](mailto:apbartok@gmail.com)**Keywords:** Gaussian approximation potential, machine learning, interatomic potential, high performance computing, ScaLAPACK, distributed linear algebra, QR decompositionOriginal Content from  
this work may be used  
under the terms of the  
[Creative Commons  
Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/).Any further distribution  
of this work must  
maintain attribution to  
the author(s) and the title  
of the work, journal  
citation and DOI.**Abstract**

We present a data-parallel software package for fitting Gaussian approximation potentials (GAPs) on multiple nodes using the ScaLAPACK library with MPI and OpenMP. Until now the maximum training set size for GAP models has been limited by the available memory on a single compute node. In our new implementation, descriptor evaluation is carried out in parallel with no communication requirement. The subsequent linear solve required to determine the model coefficients is parallelised with ScaLAPACK. Our approach scales to thousands of cores, lifting the memory limitation and also delivering substantial speedups. This development expands the applicability of the GAP approach to more complex systems as well as opening up opportunities for efficiently embedding GAP model fitting within higher-level workflows such as committee models or hyperparameter optimisation.

**1. Introduction**

Computational materials and molecular modelling have proved to be an invaluable tool in predicting new materials and processes and interpreting experimental phenomena on the microscopic level. The predictive performance of atomistic simulations strongly depends on the accuracy of the employed atomic interaction model, of which those based on solving the Schrödinger equation are generally regarded as the most reliable. However, many problems of interest remain intractable, even when using approximate solutions of the quantum mechanical problem, such as density functional theory (DFT), due to the high computational cost and its scaling with respect to system size. While interatomic potentials based on simple analytical forms open up access to significantly larger system sizes and longer simulation times, their parameterisation is often insufficiently accurate for predictive modelling. Surrogate models for quantum mechanical calculations, based on highly flexible functional forms provided by machine learning methods which are fitted using high-quality *ab initio* reference data emerged in the last two decades [1–7]. These machine learning interatomic potentials (MLIPs) reproduce the *ab initio* potential energy surface to a high accuracy in a computationally efficient way, allowing access to large time and length scale simulations [8].

In this work, we focus on the fitting aspect of MLIPs, i.e. the process that determines the model parameters based on a set of reference data points. Even though fitting is typically a one-off operation, and its computational cost leaves the cost of a subsequent simulation using the MLIP largely or completely unaffected, it can use significant resources and can be a limiting factor in applying ever increasing data bases or exploring the space of model hyperparameters. Depending on the regression method, some MLIPs are based on solving a linear system to obtain the model weights. Here we present a set of application principles that can be used to distribute the workload among multiple processes when fitting such models, allowing efficient utilisation of massively parallel computer resources. We have implemented these in the Gaussian approximation potential (GAP) framework and demonstrated excellent scaling in both memory and computational efficiency up to thousands of cores. We note that similar efforts have been made to parallelise the `fitSNAP` code used for fitting Spectral Neighbour Analysis Potential [3] linear models [9]. Other MLIP

implementations which are based on kernel methods [10] or linear regression, such as the Linear Machine Learning [11] and Atomic Cluster Expansion [4] approaches, would also benefit from similar developments.

## 2. Theory

We provide a brief overview of the GAP framework, and for more details we refer the reader to a more comprehensive review [12] of the method. GAP is a Bayesian regression method that aims to create surrogate models for the quantum mechanical interactions by using a reference database consisting of atomic configurations and associated microscopic properties obtained from *ab initio* calculations. Strategies to create such databases are discussed elsewhere [13, 14]; here we start from a set of configurations, each consisting of Cartesian coordinates with the corresponding atomic species information, and for *ab initio* data we use total energies, forces and virial stress components. As Cartesian coordinates do not transform in an invariant fashion when applying energy-conserving symmetry operations to the atomic positions, such as rotations, translations, inversion and permutation of the indices of identical atomic species, it is beneficial to first transform the Cartesian coordinates to descriptors, which form the input vectors  $\mathbf{x}_i$  of length  $d$  to the regression method.

In general, GAP approximates the total energy of a configuration  $A$  as a sum of local terms

$$E_A = \sum_{i=1}^{N_A} \varepsilon(\mathbf{x}_i), \quad (1)$$

where the local energy term  $\varepsilon$  is in the form of a sparse Gaussian process (GP) [15, 16]

$$\varepsilon(\mathbf{x}_i) = \sum_j^M c_j k(\mathbf{x}_i, \mathbf{x}_j), \quad (2)$$

where the sum is over a set of representative descriptor vectors, or sparse points  $M$ . The kernel function  $k(\mathbf{x}_i, \mathbf{x}_j)$  evaluates the similarity of descriptors  $\mathbf{x}_i$  and  $\mathbf{x}_j$ , and  $c_j$  are the regression weights that need to be fitted such that predicted properties match the *ab initio* values as closely as possible. Forces and virial stress components can be obtained by differentiating this expression with respect to atomic coordinates or lattice deformations, which is a trivial, but rather tedious operation and we omit it from here for brevity. Denoting the  $N$  *ab initio* reference properties by  $\mathbf{y}$  and predicted properties by  $\tilde{\mathbf{y}}$ , we formulate the regression problem as minimising the loss function

$$\mathcal{L} = (\mathbf{y} - \tilde{\mathbf{y}})^T \boldsymbol{\Sigma}^{-1} (\mathbf{y} - \tilde{\mathbf{y}}) + \mathbf{c}^T \mathbf{K}_{MM} \mathbf{c} \quad (3)$$

with respect to the weights  $\mathbf{c}$ . We note that elements of  $\mathbf{y}$  may be total energies or their derivatives, such as force and stress components, but here we only review the case for total energies. The matrix  $\boldsymbol{\Sigma}$  is diagonal and its elements are inversely related to the importance of each data point. While the first term is responsible for achieving a close fit to the data points, the second term is controlling overfitting via a Tikhonov regularising expression, which forces the elements of  $\mathbf{c}$  to remain small. The elements of the matrix  $\mathbf{K}_{MM}$  are built from the kernel function values  $K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$  between the sparse point set  $\{\mathbf{x}_j\}_{j=1}^M$  where we typically use  $M \ll N$ . The minimum of the loss function in equation (3) can be determined analytically as a solution of the linear system

$$(\mathbf{K}_{MM} + \mathbf{K}_{MN} \boldsymbol{\Sigma}^{-1} \mathbf{K}_{NM}) \mathbf{c} = \mathbf{K}_{MN} \boldsymbol{\Sigma}^{-1} \mathbf{y} \quad (4)$$

for  $\mathbf{c}$ , or equivalently

$$\mathbf{c} = (\mathbf{K}_{MM} + \mathbf{K}_{MN} \boldsymbol{\Sigma}^{-1} \mathbf{K}_{NM})^{-1} \mathbf{K}_{MN} \boldsymbol{\Sigma}^{-1} \mathbf{y}. \quad (5)$$

The elements of  $\mathbf{K}_{MN}$  are given by

$$K_{ij} = \sum_{\alpha=1}^{N_j} k(\mathbf{x}_i, \mathbf{x}_\alpha) \quad (6)$$

where  $\mathbf{x}_i$  is a descriptor vector from the sparse set, and  $j$  denotes a target total energy and the sum includes all descriptors that correspond to local energy terms contributing to  $y_j$ . For convenience, we use the notation  $\mathbf{K}_{MN}^T \equiv \mathbf{K}_{NM}$ . Elements of  $\mathbf{K}_{MN}$  corresponding to derivative observations are calculated similarly, using the

appropriate gradients of the kernel function  $k$ , for which further details may be found in the review article by Deringer *et al* [12].

The complexity of solving equation (4) scales with  $\mathcal{O}(M^2N)$ , which is significantly more favourable than the  $\mathcal{O}(N^3)$  scaling of a full GP implementation. However, Foster *et al* have shown [17] that the solution may lead to numerically unstable results at large data sets, i.e. uncertainties in the input lead to disproportionate errors in the output. Following their suggestions, we first define the  $(N + M) \times M$  matrix

$$\mathbf{A} = \begin{bmatrix} \Sigma^{-1/2} \mathbf{K}_{NM} \\ \mathbf{L}_{MM}^T \end{bmatrix} \quad (7)$$

where the lower triangular matrix  $\mathbf{L}_{MM}$  is the result of the Cholesky decomposition of  $\mathbf{K}_{MM}$  such that  $\mathbf{K}_{MM} = \mathbf{L}_{MM} \mathbf{L}_{MM}^T$ . Introducing  $\mathbf{b}$  by padding the vector of target properties  $\mathbf{y}$  by an  $M$ -long vector of zeros

$$\mathbf{b} = \begin{bmatrix} \mathbf{y} \\ 0 \end{bmatrix} \quad (8)$$

we rewrite equation (4) as the solution of the least-squares problem

$$\min_{\mathbf{c}} (\mathbf{A}\mathbf{c} - \mathbf{b})^T (\mathbf{A}\mathbf{c} - \mathbf{b}) \quad (9)$$

that leads to the solution in the form of

$$\mathbf{c} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b}. \quad (10)$$

A numerically stable solution can be obtained by first carrying out a QR factorisation of  $\mathbf{A} = \mathbf{Q}\mathbf{R}$  where  $\mathbf{Q}$  is orthonormal, namely, it is formed by *orthonormal* column vectors:

$$\mathbf{Q}^T \mathbf{Q} = \mathbf{I}, \quad (11)$$

while  $\mathbf{R}$  is an upper triangular matrix. Substituting the factorised form of  $\mathbf{A}$  into equation (10) results in

$$\mathbf{c} = (\mathbf{R}^T \mathbf{Q}^T \mathbf{Q} \mathbf{R})^{-1} \mathbf{R}^T \mathbf{Q}^T \mathbf{b} = \mathbf{R}^{-1} \mathbf{Q}^T \mathbf{b}. \quad (12)$$

The computational complexity of creating  $\mathbf{A}$  is determined by the cost of creating its two constituent blocks. The calculation of the upper block scales as  $\mathcal{O}(MN)$ , due to  $\Sigma$  being diagonal, while the Cholesky factorisation resulting in the lower block scales as  $\mathcal{O}(M^3)$ , resulting in an overall scaling  $\mathcal{O}(MN)$ , as  $N \gg M$ . The QR factorisation of  $\mathbf{A}$  requires  $\mathcal{O}(M^2N)$  floating point operations, hence dominating the overall cost of evaluating equation (10). We note that multiplying by  $\mathbf{R}^{-1}$  can be implemented as a series of back substitution operations, due to the upper triangular matrix form of  $\mathbf{R}$ .


### 3. Implementation

The workflow of obtaining the sparse or representative points and associated vector of weights  $\mathbf{c}$  from a set of reference *ab initio* configurations is implemented in the `gap_fit` program, and distributed as part of the software package Quantum Mechanics and Interatomic Potentials (QUIP), which is a Fortran package implementing atomistic simulation tools, including low-level functions to manipulate atomic configurations, a selection of interatomic potentials, tight-binding models and the GAP framework. The source code is publicly available on Github [18].

#### 3.1. Program structure

The `gap_fit` program is controlled via a set of command line arguments consisting of key-value pairs, which can also be passed as a configuration file. The program also requires a set of reference configurations in the extended XYZ format [19], containing any combination of total energies, forces and virial stresses, and optionally, the definition of a baseline potential model. The major steps of the fitting process are outlined in figure 1. After initialisation and reading of the command line arguments, the training structures are parsed for the number of target properties: total energies, forces and virial stress components, to determine the value of  $N$ . Based on the descriptors, the amount of storage space needed for the descriptor arrays  $\mathbf{x}$  and their derivatives  $\mathbf{x}'$  with respect to Cartesian coordinates are calculated and then allocated.

From the descriptor vectors,  $M$  are chosen as a representative (sparse) set. The procedure for choosing can be controlled by command line arguments, including selecting a random subset, clustering and



CUR-based approaches [20]. It is also possible to provide the chosen representative points via files, an option we make use of for the parallel version (see section 3.2).

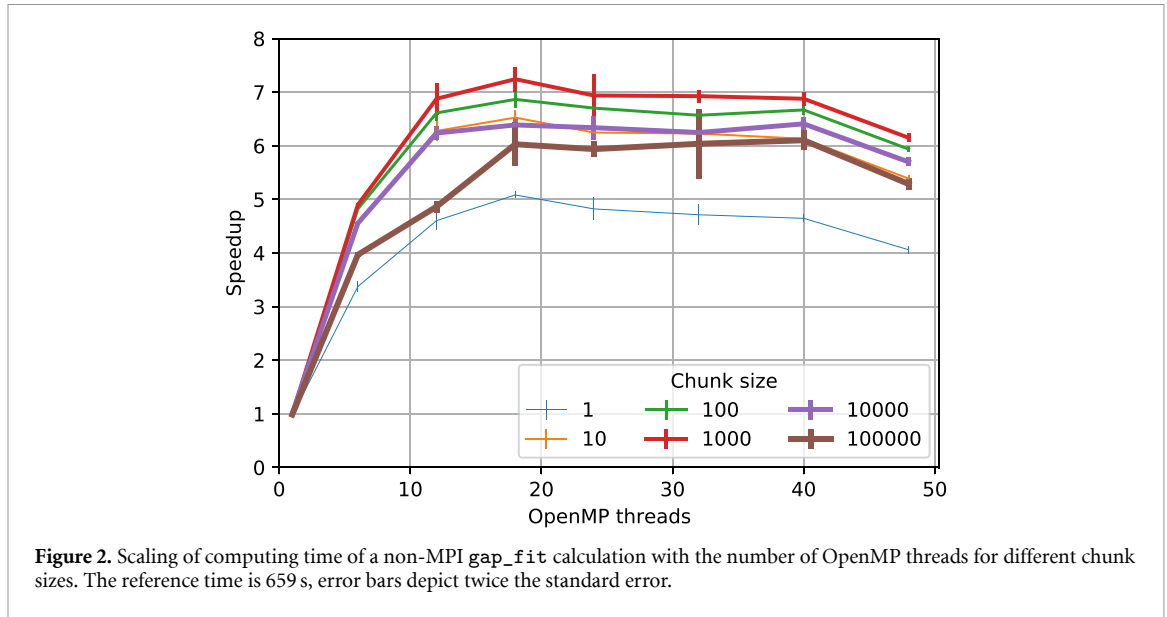
After setting the sparse points, the covariance matrices  $\mathbf{K}_{MN}$  and  $\mathbf{K}_{MM}$  are calculated. From these, matrix  $\mathbf{A}$  is constructed and equation (10) is solved via QR decomposition using linear algebra routines from the Linear Algebra Package (LAPACK) library for single node applications.

The intermediate processing, such as the computation of the elements of covariance matrices had already been augmented by Open Multiprocessing (OpenMP) directives along the target data dimension  $N$ , which leads to a thread-based parallelisation on a single process. This, however, restricts the program to the memory and processing resources of a single node, and performance is further limited by the fact that the speed a computational core can access an arbitrary memory region is inhomogeneous due to the physical layout of the memory. That results in a decrease of efficiency when additional threads are employed, leading to a degradation of performance which prevents full utilisation of all available cores in a node. We present the parallel scalability of a test problem in figure 2, where we varied the size of contiguous subsets of OpenMP loops, referred to as *chunks*. These calculations were run on the Avon cluster at the University of Warwick (specifications: Intel Xeon Platinum 8268 (Cascade Lake) processors (2.9 GHz), 48 Cores).

As an example of the limitations imposed by the OpenMP implementation of `gap_fit`, the practical problem of fitting a GAP for carbon [21]—one of the largest single-element training datasets assembled to date—took more than 6 days on a single node and required more than 1 TB memory to accommodate the design and covariance matrices [22]. This restricted the ability of practitioners to build complete training sets or to experiment with choices of hyperparameters.

### 3.2. Multi-process parallelisation

To go beyond the limitations posed by the shared memory requirement, poor parallel performance, and specialist hardware, we propose a multi-process framework with data distribution and inter-node communication. We have established in section 2 that both the memory requirement and the computational effort scale linearly with the number of target properties  $N$ , therefore it is convenient to distribute memory and tasks along this dimension of the problem.



The two most memory intensive operations are the calculation of the descriptor vectors together with their gradients, and the covariance matrices. The ratio of these depends strongly on the particulars of the fitting problem, especially the dimensionality  $d$  of descriptor vectors and the number of sparse points  $M$ . In our parallel scheme, we distribute atomic configurations across independent processes, such that the number of target properties are as even as possible. We note that the size of individual atomic configurations may be highly inhomogeneous, therefore the number of forces per configuration can vary substantially across the database, necessitating an extra step that determines the optimal spread of data. We have employed an algorithm that first collects configurations in a list and sorts them by descending number of target properties. We then assign the largest (by target property) unassigned configuration to the process which currently has the least total number of target properties. This process repeats until the list is exhausted.

With the configurations allotted to Message Passing Interface (MPI) processes, the descriptor calculations may proceed locally, and once completed, individual portions of  $\mathbf{K}_{MN}$ , denoted as  $\mathbf{K}_{Mn}$  can be evaluated. For this, local copies of the sparse set of  $M$  descriptor values need to be present locally, the particulars of which we discuss later in section 3.3. The naïve solution of the linear system represented by equation (4) may be adapted trivially to a distributed  $\mathbf{K}_{MN}$ : the terms  $\mathbf{K}_{MN}\Sigma^{-1}\mathbf{K}_{NM}$  and  $\mathbf{K}_{MN}\Sigma^{-1}\mathbf{y}$  can be calculated locally and reduced across processes as

$$\mathbf{K}_{MN}\Sigma^{-1}\mathbf{K}_{NM} = \sum_{n \in N} \mathbf{K}_{Mn}\Sigma_n^{-1}\mathbf{K}_{nM} \quad (13)$$

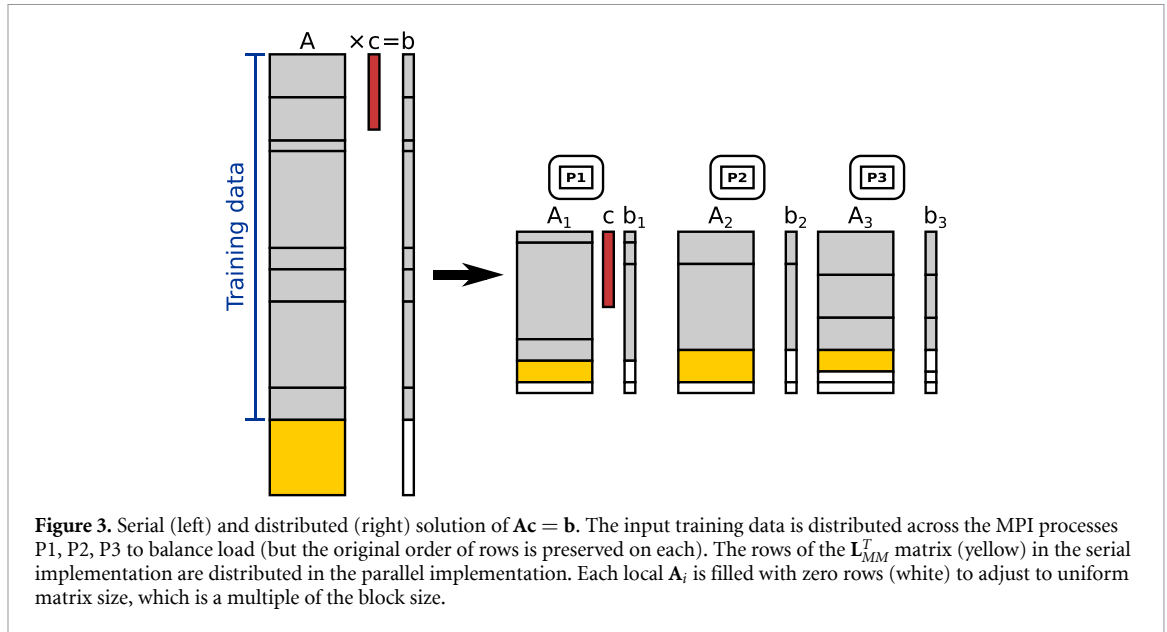
and

$$\mathbf{K}_{MN}\Sigma^{-1}\mathbf{y} = \sum_{n \in N} \mathbf{K}_{Mn}\Sigma_n^{-1}\mathbf{y}_n \quad (14)$$

where we denote distributed blocks of  $\Sigma$  and  $\mathbf{y}$  by  $\Sigma_n$  and  $\mathbf{y}_n$ , respectively. The rest of the calculation only involves matrices up the size of  $M \times M$ . However, the direct solution, as described in section 2 is numerically unstable, therefore we need to adapt the solution based on the QR-factorisation.

The Scalable Linear Algebra Package (ScaLAPACK) library provides some of the linear algebra features of the LAPACK library for distributed matrices, most commonly leveraging the MPI framework for communication between nodes, which is widely available on computing clusters. We chose to leverage the ScaLAPACK implementation, therefore we need to take ScaLAPACK's data distributing principles in consideration. The procedure names are the same as for LAPACK but with a prefix `p`, e.g. the QR-factorisation subroutine is `pdgeqrf` instead of `dgeqrf`. For the rest of our discussion, we will use the prefixed names.

ScaLAPACK asserts that matrices are block-cyclicly distributed on a 2D processor grid. This is a generalisation of cyclic and block distribution, both of which can be used as special cases. Considering a matrix  $\mathbf{A}_{R \times C}$  with  $R$  rows and  $C$  columns, we can cut it into blocks  $\mathbf{a}_{r \times c}$ . The last blocks in each row or column may have fewer columns or rows. The blocks are then distributed in a round-robin fashion amongst the processors in the  $p \times q$  processors grid, wrapping around the grid until all blocks have been assigned.



For our use-case we start by considering a block distribution along the rows for a processor grid of  $p \times 1$ . This entails a row block size equal to the local number of rows for each process ( $\mathbf{a}_{r \times C}$  and  $\mathbf{b}_{r \times 1}$  with  $r = \lceil R/p \rceil$ ). We fill these blocks by assigning each structure (i.e. several rows of  $\mathbf{a}$ ) to a single process, thereby each atomic configuration is local on exactly one process. The solution to  $\mathbf{Ac} = \mathbf{b}$  is invariant to swapping rows of  $\mathbf{A}$  as long as the corresponding entries in  $\mathbf{b}$  are swapped accordingly. This allows us to choose the row block size freely while arriving at the same result irrespective of the assignment of atomic configurations to processes. The column block size is unrestricted, since each row is fully assigned to a single process.

Our algorithm, as described above and presented in figure 3, distributes atomic configurations and rows of  $\mathbf{L}_{MM}^T$  such that each local  $\mathbf{A}_i$  block is as equal in size as possible. ScaLAPACK requires that all processes use a uniform block factor for all their blocks. To fill the gaps left by the distribution a padding of zero rows (i.e. rows filled with zeroes) is inserted into both  $\mathbf{A}$  and  $\mathbf{b}$ . The distribution strategy and the block size settings of ScaLAPACK should ensure that the number of padding rows are kept to a minimum to prevent afflicting memory and efficiency penalties.

Solving the linear system via QR decomposition with ScaLAPACK is split into three steps. First,  $\mathbf{A}$  is converted using `pdgeqrf` into the upper triangular matrix  $\mathbf{R}$  and the Householder reflectors  $\mathbf{H}$ , which occupy the remaining lower triangular elements of  $\mathbf{A}$ . The latter is accompanied by an array  $\boldsymbol{\tau}$  of weights. Reflectors and weights are then used by `pdormqr` to perform the multiplication  $\mathbf{Q}^T \mathbf{b}$ . Finally, the linear system represented by the upper triangle matrix  $\mathbf{R}$  is solved by `pdtrtrs`, utilising the backsubstitution algorithm, to give  $\mathbf{c} = \mathbf{R}^{-1} \mathbf{Q}^T \mathbf{b}$ .

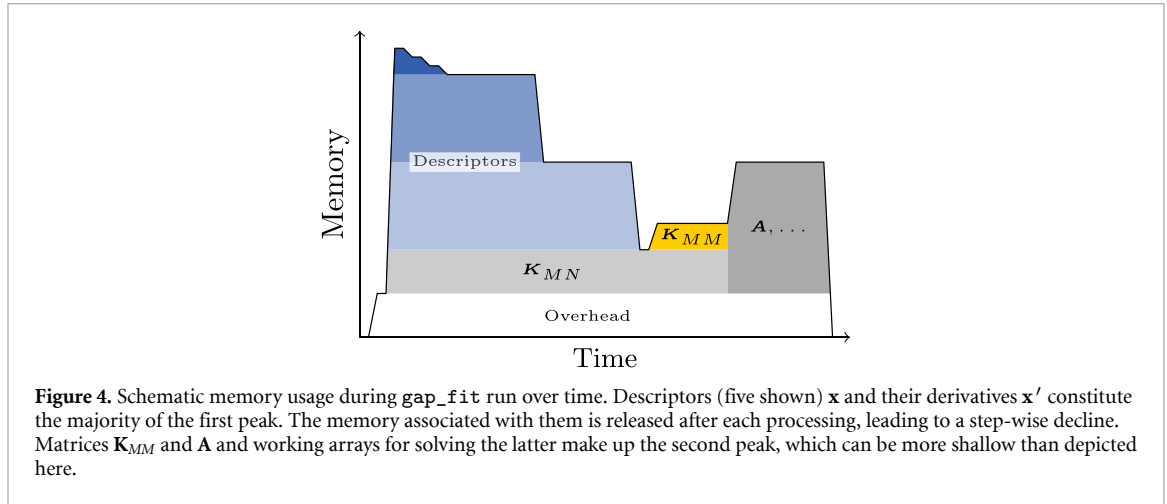
We note that there is a requirement in `pdtrtrs` that the row and column block sizes must be equal. Setting the column block size ( $c$ ) to the generally much larger row block size ( $r$ ) is formally possible, but this drastically increases the size of the working arrays the ScaLAPACK routines require, which scale with the square of the column block size ( $\propto c^2 + rc$ ). Setting instead the row block size ( $r$ ) to the column block size ( $c$ ) implies adding additional zero rows for padding the local matrices to maintain the divisibility by the block factor and thus the assignment of the configurations to the processes. Both of these approaches result in increased memory requirements and a deterioration of computational efficiency.

However, it is possible to exploit the fact that our distribution strategy relies on a single processor column ( $q = 1$ ), so changing the column block size does not affect the distribution of the data. We can therefore use one column block size for the first two calls (`pdgeqrf`, `pdormqr`) and then change that value for the third call (`pdtrtrs`) to fulfill its requirement without decreasing the efficiency of the former calls.

Being able to control both block sizes independently revealed that a moderate column block size of about 100 is optimal for both memory usage and efficiency. For such a setting, the row block size does not have a significant impact on parallel efficiency.

### 3.3. Sparse point selection

In `gap_fit`, the set of  $M$  sparse points are typically determined as a subset of all descriptor values, although for low-dimensional descriptors such as bond length it is convenient to use a uniform grid. Depending on the method, the selection of sparse points may depend on the values of descriptor vectors calculated from the



entire training data set. If the descriptors are distributed, clustering or CUR-based methods require fine-tuned communication between the processes and for simplicity, we suggest a two-step workflow. Since the calculation of descriptor vectors—without their gradients—is not computationally expensive and memory storage is not a concern, sparse point selection can be performed using serial execution.

We first run `gap_fit` on a single process, optionally using OpenMP threading, to select sparse points which are written into files, and terminating the run, which can be achieved by the command line argument `sparsify_only_no_fit=T`. The output is converted to input via a helper script for the subsequent run using MPI processes. This step can be skipped if the sparse points file has been provided by external means or can be reused from an earlier calculation.

### 3.4. Peak memory usage

One of the pitfalls of running a formerly serial program in parallel with distributed data is that duplicate data may accumulate unnecessarily, especially if multiple processes are run on the same node, and therefore shared memory can be utilised. For example, it is convenient to calculate the matrix  $\mathbf{K}_{MM}$  on each process because it only depends on the sparse points, and its size does not depend on the training set. However, each process requires only a small part of the resulting matrix  $\mathbf{L}_{MM}$ , and storing multiple copies of  $\mathbf{K}_{MM}$  adds an unnecessary constant overhead to the memory requirements. To prevent the allocation of possibly several GB memory per process,  $\mathbf{K}_{MM}$  is only calculated on a single process, then converted to  $\mathbf{L}_{MM}$ , and only the necessary rows are distributed via `mpi_scatterv` calls to independent MPI processes.

It is also important to avoid inadvertent duplication of data when converting between data structures used by different parts of the program. This can be alleviated by performing calculations directly on the data memory as LAPACK does. For user-defined types we use pointers to the original matrices to prevent copying of data. Further, source data of relevant size is deallocated after final usage. This decreases the memory overhead per MPI process and therefore also the peak memory requirement of the program.

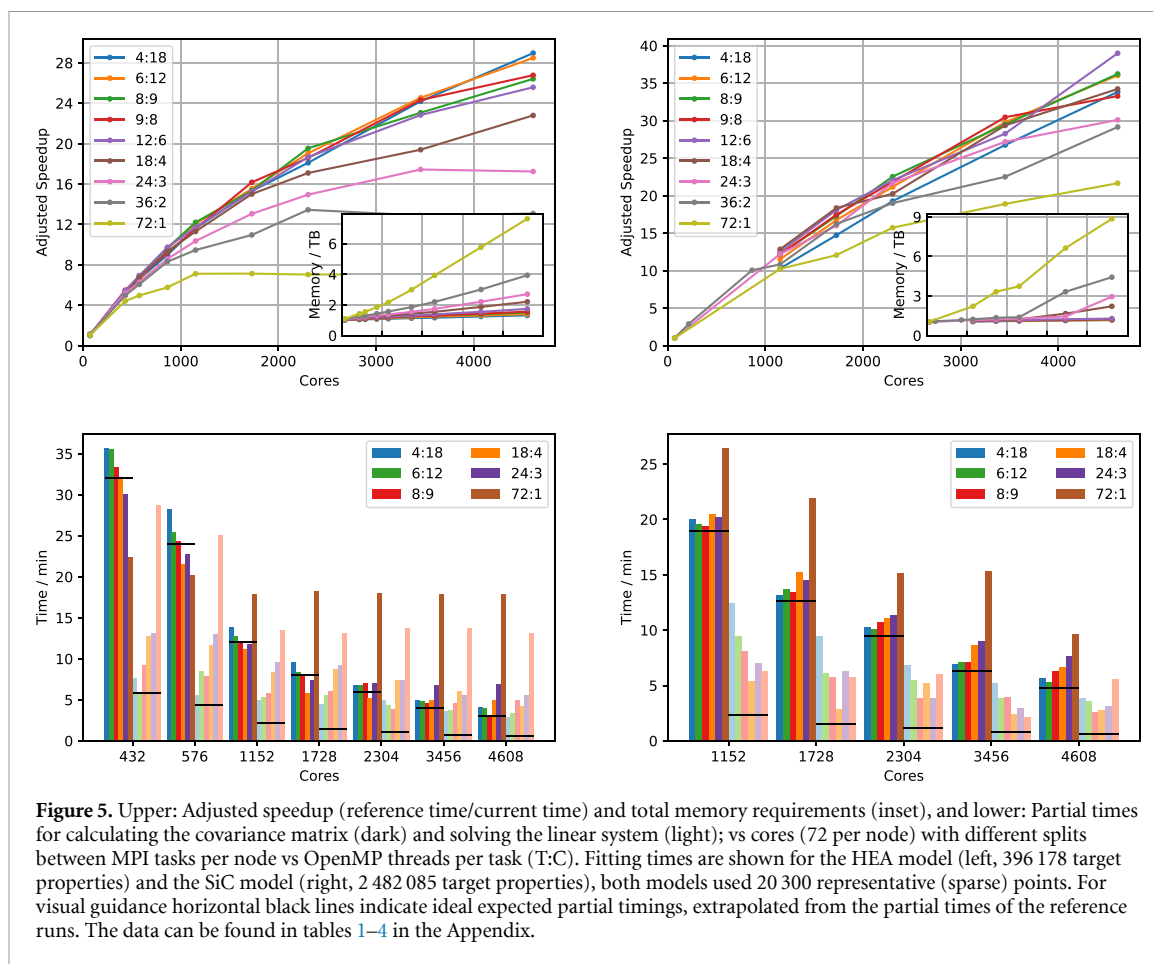
Figure 4 shows schematically how the memory usage of `gap_fit` run changes over time. For our program there are two parts of the execution which may lead to peak memory usage. The main one is after allocating the descriptors, especially the derivatives  $\mathbf{x}'$ . After the covariance calculation of each descriptor, we deallocate the corresponding source data. This is reflected by the step-wise decline of the memory.

The other peak manifests towards the end of a program run when the matrices  $\mathbf{K}_{MM}$  and then  $\mathbf{A}$  are assembled and the linear system is subsequently solved. ScaLAPACK requires additional work arrays for some of its routines depending on the block sizes of the distributed matrices, especially the column block size.

## 4. Practical examples

Initial proof-of-principle fitting runs of a silicon dataset [23] showed that the fitting weights from an MPI run are comparable to those from serial runs when using the same representative set. The difference can be attributed to numerical uncertainties noting that even two runs initiated with identical parameters may differ to the same magnitude. The difference can be attributed to the fact that the order of numerical operations is non-deterministic and the floating point arithmetic is neither associative nor distributive, leading to small differences in covariance matrices in different executions. Ill-conditioning of matrix  $\mathbf{A}$  amplifies the noise due to the different order of operations, leading to only a few comparable significant digits in the resulting weights. We have therefore tested the accuracy of the predictions with the resulting potential models using





the  $\Delta$  metric suggested to compare DFT packages [24]. We found that equivalent GAP models differ only up to  $1 \mu\text{eV}$ , indicating excellent reproducibility despite the differences in the weights.

We then applied the implementation with varying proportions of MPI processes and OpenMP threads on two example training sets, consisting of an high-entropy alloy (HEA) and silicon carbide (SiC) datasets.

These calculations were performed on the Raven cluster of the Max Planck Computing and Data Facility. Each node has 72 cores and either 256 GB or 512 GB of RAM. For single-node calculations we used high-memory nodes with 2048 GB of RAM. The specifications are: Intel Xeon IceLake-SP processors (2.4 GHz), 72 Cores (2 NUMA domains), Mellanox HDR InfiniBand network (100 Gbit/s, pruned fat-tree topology) [25].

In both cases we combined two-body descriptors with Smooth Overlap of Atomic Positions (SOAP) descriptors [26]. We assigned separate two-body descriptors for each pair of species, and separate SOAP descriptors for each species.

The upper part of figure 5 depicts the inverse time relation with respect to the number of cores for different MPI to OpenMP ratios (T:C), e.g. ‘36:2’ means that 36 MPI processes were used per node (with 72 cores), each with two OpenMP threads. This resembles Amdahl’s law

$$S(n) = t(1)/t(n) = 1/[(1 - p) + p/n], \quad (15)$$

where the speedup  $S$  describes the relative time for a serial run  $t(1)$  versus a parallelised one  $t(n)$ . It depends on the number of cores  $n$  and the relative time  $p$  spent in portions that benefit from multiple cores. Our training systems were too large to be run on a single core within the maximum wall-time limit, so we used the highest time available instead for our adjusted speedup ( $S^*$ ). Because of this, the values are only comparable within the same training set. Note that these timings may contain some random noise due to other calculations on the cluster (albeit not on the same nodes) and generally unpredictable behavior of interconnect, network, hardware or the operating system in general. The insets show the total memory required for these calculations across all nodes, estimated from the average resident set size (AveRSS) as queried from the queueing system.

The lower part of figure 5 shows the absolute times for the calculation of the covariance matrix, and for preparing and solving the linear system. As a visual guide horizontal lines indicate the ideal time extrapolated

from the respective partial times of the (longest) reference calculation. The remaining time to the total run time is small and not depicted for readability.

#### 4.1. High-entropy alloy

The high-entropy MoNbTaVW alloy (HEA) training set [27] consists of 2329 configurations—each containing an uneven number of atoms—with 2329 total energies, 383 739 forces, and 10 110 virials for a total of  $N = 396\,178$  target properties. With 20 sparse points per two-body descriptor (15) and 4000 per SOAP descriptor (5) the total number of sparse points is  $M = 20\,300$ . Thus,  $\mathbf{A}$  consists of  $n_{\mathbf{A}} = 8\,454\,503\,400$  elements and occupies about 67.6 GB.

Looking at figure 5, the speedups for ratios between 4:18 and 12:6 do not differ much in our timings. Calculations with more MPI processes lose efficiency with increasing number of cores.

For example, omitting OpenMP (72:1) decreases the speedup gain around 500 cores, staying at constant time between 864 and 1152 cores. The same happens for 36:2 ratio between 1728 and 2304 cores, and for 24:3 between 2304 and 3456 cores. This behaviour stems from the choice of our implementation: it splits the training set along the structures, which cannot be done arbitrarily for a finite set. For these 2329 configurations, the limit is at about two structures per MPI process.

The reference time for this set is 13 815 s (3.8 h), obtained from the single (high-memory) node calculation with a 4:18 split of cores. Other ratios improve only a little on top of that, the highest being 12:6, 24:3, and 36:2 with a speedup of 1.13 each. The 72:1 run does not achieve that (1.06). Using six nodes (432 cores) increases the speedup to between 5.49 (9:8) and 4.42 (72:1), which translates to relative increases between 5.09 and 4.17 compared to a single node. As mentioned before, the benefits dwindle fast for the 72:1 split from factor 1.31 (total 5.78) despite doubling (twelve nodes) to the constant speedup around 16 nodes of ca. 7.10. The 36:2 ratio achieves a speedup of 13.45 at its best but decreases to about 13.0 beyond that. The same happens, less pronounced, for 24:3, from a maximum of 17.44 at 48 cores down to 17.25. The highest speedup in this set is 28.96 (4:18 on 64 nodes). The most efficient w.r.t. to the used resources is 18:4 on a single node.

Using more MPI processes also comes at the cost of a memory overhead, which increases approximately linearly with the number of cores. The higher the portion of MPI usage, the steeper this overhead is. In fact, the graphs coincide if they are plotted against the number of MPI processes (not shown): in that case the slope ranges between 0.9 and 1.5 GB per MPI process. For 72:1 this results in 2.17 TB on 16 nodes and 7.63 TB on 64. The former is comparable to the 2.20 GB 18:4 uses on 64 nodes, since both apply 1152 MPI processes.

Looking at the partial timings, solving the covariance is close to the ideal performance when increasing the number of nodes. The reference times are 3.2 h for covariance and 0.6 h for the QR preparation and solve from the 4:18 ratio on one node. Calculations with a higher number of MPI tasks per node needs less time for a smaller number of nodes up to the point of saturation when this time stays constant. Solving of the linear system takes less time than the covariance but does not benefit as much from an increasing number of nodes. This is to be expected as more inter-node communication slows down the computation, while no communication is needed for the covariance. Using all cores for MPI results in a severe penalty on the solving time. The remaining runtime is negligible (up to 3 min).

#### 4.2. Silicon carbide

The 4865 silicon carbide (SiC) systems of this training set contain 4865 energies, 2448 030 forces, 29 190 virials for a total of  $N = 2\,482\,085$  target properties. With 100 sparse points per two-body descriptor (3) and 10 000 per SOAP descriptor (2) the total number of sparse points is  $M = 20\,300$ . Thus,  $\mathbf{A}$  consist of  $n_{\mathbf{A}} = 50\,798\,415\,500$  elements and occupies 406.4 GB.

Due to the much larger training set, not all node configurations from the HEA set were viable, especially for lower node numbers. The reference time is 20 472 s (5.7 h) for a 24:3 run on a single node. The trend of an equal split between MPI and OpenMP being the preferred choice also holds for this set but the processes are not saturated as rapidly as in the case of the HEA systems. This effect may start between 3456 and 4608 cores for the full MPI run (72:1), which is later than in the HEA set even taking the structure numbers into account because of the proportionally lower number of small structures in this larger set of only two species. The highest speedup of 38.99 is achieved by 12:6 on 64 nodes, similarly the 18:4 split reaches the best value of 12.80 on 16 nodes whereas 36:2 gives 10.84 and 4:18 only 10.31.

Our method of querying the queueing system underestimates the average memory usage for some of the calculations. The general trend is, however, reliable to the HEA case, as shown by a similar values for high core numbers. The memory overhead per MPI process is nominally between 0.3 and 1.8 GB. On 16 nodes the memory usage for 72:1 is given as 2.22 TB, on 64 nodes as 8.85 TB.

For the partial timings we have similar trends as in the HEA case. The reference times are 5.1 h for covariance and 0.6 h for the QR preparation and solve from the 24:3 ratio on one node. Strikingly, the 72:1 covariance times are much higher than for other splits. The other times are quite similar between different ratios, in agreement with the total speedup. The remaining runtime takes up to 1 min.

## 5. Conclusion and outlook

The recent addition of MPI parallelisation to our program `gap_fit` by using the ScaLAPACK library makes it possible to split the training data evenly into batches to be processed independently up to the solving of the linear system. It alleviates the need for high-memory nodes so commodity HPC nodes may be used for arbitrarily large fitting problems. The efficiency of the MPI parallelisation is comparable to that of OpenMP for all parallelisation strategies except fully MPI, effectively extending the previous efficiency to multiple nodes. Thus larger training sets do not impede the computation and more sparse points can be used, increasing the accuracy of the model.

We showed the time scaling and memory requirements for varying proportions of MPI processes vs OpenMP threads in two example training sets, consisting of an high-entropy alloy (HEA) and silicon carbide (SiC). It is generally advisable to use at least some cores on each node for MPI processes, so even on a single node benefits from this new feature. It is especially effective for larger training sets while the sparse points are covered by the OpenMP threads. One should keep the total number of MPI processes below some fraction of the total number of structures, e.g. 0.5 so that an even distribution is still possible. Analysing the partial times of calculating the covariance matrix and solving the linear system confirmed that the benefit stems from the almost perfectly parallel covariance part due to the splitting of configurations. The inter-node communication of the ScaLAPACK routines prevents the solving part to benefit as much from a larger number of nodes.

The memory overhead due to the parallelisation has been reduced but is still significant. Depending on the available memory resources, a higher share of OpenMP threads is preferable, coinciding with the conclusion from the efficiency evaluation. An even smaller memory footprint may be achieved in a further development, albeit the impact is not as severe since the recommended number of MPI processes is moderate.

The highest impact on both MPI and non-MPI memory requirements would be to fully restructure the descriptor processing loop so that each descriptor is processed fully before the next one.

In practical tests we have seen that the parallel `gap_fit` code can decrease the time required to fit potentials from days to tens of minutes. We anticipate this will be an important step to enable potential fitting to be embedded within other higher-level workflows such as active learning [13, 14], committee models [28] as well as enabling model hyperparameters to be tuned or optimised, known to be important for improved uncertainty quantification [14].

## Data availability statement

The data that support the findings of this study are available upon reasonable request from the authors.

## Acknowledgments

We thank Harry Tunstall for providing the SiC dataset and early testing as well as Gábor Csányi and Miguel Caro for useful discussions. This work was financially supported by the NOMAD Centre of Excellence (European Commission grant agreement ID 951786) and the Leverhulme Trust Research Project Grant (RPG-2017-191). ABP acknowledges support from the CASTEP-USER project, funded by the Engineering and Physical Sciences Research Council under the grant agreement EP/W030438/1. We acknowledge computational resources provided by the Max Planck Computing and Data Facility provided through the NOMAD CoE, the Scientific Computing Research Technology Platform of the University of Warwick, the EPSRC-funded HPC Midlands+ consortium (EP/T022108/1) and ARCHER2 ([www.archer2.ac.uk/](http://www.archer2.ac.uk/)) via the UK Car-Parinello consortium (EP/X035891/1). We thank the technical staff at each of these HPC centres for their support.

For the purpose of open access, the author has applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising from this submission.

## Appendix. `gap_fit` options and timings

The following input is split into separate entries here for readability. It needs to be expanded for use.

For the HEA set:

```
at_file=data.xyz gap=GAP default_sigma={0.005 0.2 0.05 0.0}
energy_parameter_name=dft_energy force_parameter_name=dft_forces
sparse_jitter=1.0E-8 gp_file=gap.xml rnd_seed=999
GAP={SUBGAP1:SUBGAP2}
SUBGAP1={distance_2b cutoff=4.5 delta=10.0 covariance_type=ard_se
theta_uniform=0.75 n_sparse=20 sparse_method=uniform}
SUBGAP2={soap l_max=4 n_max=8 atom_sigma=0.5 zeta=2 cutoff=3.5
cutoff_transition_width=0.5 central_weight=1.0 n_sparse=4000 delta=0.1
covariance_type=dot_product sparse_method=cur_points}
```

For the SiC set: rnd\_seed=84616238, n\_sparse=100, and n\_sparse=10000, respectively.

**Table 1.** Time  $t$  (in s), partial times  $t_{cov}$  for covariance matrix and  $t_{QR}$  for the QR solve, adjusted speedup  $S^*$ , and estimated memory (in TB) in the HEA training set. N: nodes, T: MPI tasks per node, C: OpenMP threads per task.

N	T	C	$t_{cov}$	$t_{QR}$	$t$	$S^*$	Mem
1	4	18	11545.54	2090.22	13815	1.00	1.02
1	6	12	11461.97	1450.42	13052	1.06	1.02
1	8	9	11580.61	998.30	12699	1.09	1.03
1	9	8	11324.90	1316.42	12757	1.08	1.03
1	12	6	11515.25	756.68	12376	1.12	1.03
1	18	4	11042.48	1072.07	12208	1.13	1.04
1	24	3	11613.63	562.52	12262	1.13	1.04
1	36	2	11571.76	596.12	12248	1.13	1.05
1	72	1	11128.19	1798.22	13002	1.06	1.09
6	4	18	2136.17	454.57	2706	5.11	1.05
6	6	12	2134.75	338.10	2540	5.44	1.06
6	8	9	2003.19	549.59	2618	5.28	1.07
6	9	8	2019.54	431.06	2516	5.49	1.08
6	12	6	1955.56	512.57	2529	5.46	1.10
6	18	4	1910.88	762.27	2734	5.05	1.14
6	24	3	1804.70	789.55	2653	5.21	1.18
6	36	2	1753.38	963.17	2775	4.98	1.26
6	72	1	1341.73	1724.40	3126	4.42	1.42
8	4	18	1692.67	330.79	2093	6.60	1.06
8	6	12	1522.45	506.49	2098	6.58	1.08
8	8	9	1457.35	472.29	1993	6.93	1.09
8	9	8	1458.57	506.79	2028	6.81	1.10
8	12	6	1423.03	511.80	1995	6.92	1.13
8	18	4	1293.30	696.93	2048	6.75	1.18
8	24	3	1360.67	777.20	2196	6.29	1.22
8	36	2	1287.31	929.63	2274	6.08	1.30
8	72	1	1211.06	1506.51	2779	4.97	1.54
12	4	18	1094.36	377.24	1536	8.99	1.08
12	6	12	1043.46	331.62	1440	9.59	1.10
12	8	9	1004.06	366.18	1428	9.67	1.13
12	9	8	974.53	467.68	1501	9.20	1.15
12	12	6	915.00	444.80	1417	9.75	1.19
12	18	4	876.25	546.52	1480	9.33	1.26
12	24	3	891.94	660.54	1609	8.59	1.30
12	36	2	696.66	908.88	1661	8.32	1.42
12	72	1	1076.56	1255.26	2392	5.78	1.85
16	4	18	831.71	293.98	1187	11.64	1.10
16	6	12	766.59	320.15	1145	12.07	1.13
16	8	9	731.07	343.83	1131	12.21	1.17
16	9	8	695.14	438.43	1192	11.59	1.19
16	12	6	696.03	425.87	1178	11.73	1.22
16	18	4	666.86	499.93	1222	11.31	1.30
16	24	3	702.33	574.44	1333	10.36	1.38
16	36	2	619.84	782.44	1459	9.47	1.57
16	72	1	1070.94	807.78	1939	7.12	2.17
24	4	18	575.69	268.59	902	15.32	1.14
24	6	12	504.06	329.16	889	15.54	1.19
24	8	9	485.88	358.45	900	15.35	1.24
24	9	8	478.54	319.92	854	16.18	1.27
24	12	6	469.10	381.74	905	15.27	1.30
24	18	4	344.33	519.75	920	15.02	1.43
24	24	3	445.16	555.32	1058	13.06	1.55
24	36	2	555.51	647.66	1260	10.96	1.86
24	72	1	1090.61	783.59	1936	7.14	2.99
32	4	18	407.57	299.81	763	18.11	1.17
32	6	12	407.41	262.24	725	19.06	1.24
32	8	9	421.34	231.88	707	19.54	1.31
32	9	8	374.51	309.80	743	18.59	1.30
32	12	6	373.78	311.40	740	18.67	1.38

(Continued.)

Table 1. (Continued.)

$N$	$T$	$C$	$t_{cov}$	$t_{QR}$	$t$	$S^*$	Mem
32	18	4	310.40	443.29	808	17.10	1.55
32	24	3	420.02	444.83	924	14.95	1.74
32	36	2	557.74	414.69	1027	13.45	2.20
32	72	1	1081.61	820.29	1961	7.04	3.93
48	4	18	298.81	216.20	571	24.19	1.25
48	6	12	289.29	219.85	563	24.54	1.31
48	8	9	273.59	270.72	599	23.06	1.39
48	9	8	213.21	299.95	568	24.32	1.43
48	12	6	228.32	322.18	605	22.83	1.55
48	18	4	295.86	359.62	712	19.40	1.87
48	24	3	403.50	332.40	792	17.44	2.21
48	36	2	548.22	467.59	1072	12.89	3.01
48	72	1	1067.96	825.62	1952	7.08	5.77
64	4	18	248.17	171.94	477	28.96	1.32
64	6	12	234.53	197.51	485	28.48	1.39
64	8	9	175.12	293.97	523	26.41	1.51
64	9	8	177.17	283.84	516	26.77	1.56
64	12	6	230.36	255.08	540	25.58	1.74
64	18	4	296.01	254.14	606	22.80	2.20
64	24	3	411.00	335.30	801	17.25	2.70
64	36	2	545.20	454.38	1056	13.08	3.94
64	72	1	1072.27	784.67	1916	7.21	7.63

Table 2. Pivot table of the adjusted speedups for the HEA training set.

$t$	$T$								
	4	6	8	9	12	18	24	36	72
1	1.00	1.06	1.09	1.08	1.12	1.13	1.13	1.13	1.06
6	5.11	5.44	5.28	5.49	5.46	5.05	5.21	4.98	4.42
8	6.60	6.58	6.93	6.81	6.92	6.75	6.29	6.08	4.97
12	8.99	9.59	9.67	9.20	9.75	9.33	8.59	8.32	5.78
16	11.64	12.07	12.21	11.59	11.73	11.31	10.36	9.47	7.12
24	15.32	15.54	15.35	16.18	15.27	15.02	13.06	10.96	7.14
32	18.11	19.06	19.54	18.59	18.67	17.10	14.95	13.45	7.04
48	24.19	24.54	23.06	24.32	22.83	19.40	17.44	12.89	7.08
64	28.96	28.48	26.41	26.77	25.58	22.80	17.25	13.08	7.21

**Table 3.** Time  $t$  (in s), partial times  $t_{cov}$  for covariance matrix and  $t_{QR}$  for the QR solve, adjusted speedup  $S^*$ , and estimated memory (in TB) in the SiC training set. N: nodes, T: MPI tasks per node, C: OpenMP threads per task.

$N$	$T$	$C$	$t_{cov}$	$t_{QR}$	$t$	$S^*$	Mem
1	24	3	18185.23	2232.89	20472	1.00	1.03
1	36	2	18124.19	1736.61	19911	1.03	1.04
1	72	1	18551.53	1667.12	20265	1.01	1.05
3	36	2	6247.50	780.98	7067	2.90	1.08
12	36	2	1796.59	200.76	2035	10.06	1.18
16	4	18	1198.94	746.59	1986	10.31	1.06
16	6	12	1175.45	565.38	1776	11.53	1.08
16	8	9	1160.58	485.56	1681	12.18	1.10
16	9	8	1198.35	442.01	1677	12.21	1.11
16	12	6	1162.79	437.15	1634	12.53	1.13
16	18	4	1229.67	322.71	1592	12.86	1.14
16	24	3	1212.87	422.35	1674	12.23	1.17
16	36	2	1317.43	532.71	1889	10.84	1.23
16	72	1	1583.19	374.91	1998	10.25	2.22
24	4	18	785.98	568.33	1390	14.73	1.09
24	6	12	820.59	367.30	1222	16.75	1.11
24	8	9	806.53	345.81	1184	17.29	1.14
24	9	8	837.33	300.58	1172	17.47	1.15
24	12	6	824.66	274.41	1131	18.10	1.14
24	18	4	913.24	168.86	1115	18.36	1.18
24	24	3	869.64	375.20	1279	16.01	1.23
24	36	2	1058.73	168.08	1261	16.23	1.35
24	72	1	1315.80	342.81	1695	12.08	3.32
32	4	18	614.76	411.59	1062	19.28	1.10
32	6	12	607.02	325.76	967	21.17	1.14
32	8	9	640.56	232.78	908	22.55	1.17
32	9	8	627.46	275.51	937	21.85	1.15
32	12	6	602.28	291.93	928	22.06	1.17
32	18	4	665.32	312.82	1011	20.25	1.23
32	24	3	682.36	228.29	944	21.69	1.28
32	36	2	802.91	241.16	1077	19.01	1.39
32	72	1	908.52	358.63	1301	15.74	3.74
48	4	18	413.15	314.09	765	26.76	1.14
48	6	12	424.74	228.41	687	29.80	1.15
48	8	9	423.08	237.51	694	29.50	1.18
48	9	8	471.31	166.77	672	30.46	1.19
48	12	6	446.73	244.14	724	28.28	1.24
48	18	4	519.34	142.77	697	29.37	1.66
48	24	3	540.43	178.59	753	27.19	1.45
48	36	2	653.71	220.68	909	22.52	3.32
48	72	1	917.48	129.88	1083	18.90	6.63
64	4	18	340.06	230.36	605	33.84	1.18
64	6	12	319.72	214.31	568	36.04	1.18
64	8	9	376.70	155.10	565	36.23	1.22
64	9	8	353.46	228.42	615	33.29	1.24
64	12	6	350.17	141.57	525	38.99	1.29
64	18	4	395.85	167.28	598	34.23	2.22
64	24	3	456.60	189.24	680	30.11	2.95
64	36	2	447.39	217.53	702	29.16	4.43
64	72	1	576.34	331.77	945	21.66	8.85

**Table 4.** Pivot table of the adjusted speedups for the SiC training set. N: nodes, T: MPI processes per node.

N	T								
	4	6	8	9	12	18	24	36	72
1							1.00	1.03	1.01
3								2.90	
12								10.06	
16	10.31	11.53	12.18	12.21	12.53	12.86	12.23	10.84	10.25
24	14.73	16.75	17.29	17.47	18.10	18.36	16.01	16.23	12.08
32	19.28	21.17	22.55	21.85	22.06	20.25	21.69	19.01	15.74
48	26.76	29.80	29.50	30.46	28.28	29.37	27.19	22.52	18.90
64	33.84	36.04	36.23	33.29	38.99	34.23	30.11	29.16	21.66

## ORCID iDs

Sascha Klawohn  <https://orcid.org/0000-0003-4850-776X>

James R Kermode  <https://orcid.org/0000-0001-6755-6271>

Albert P Bartók  <https://orcid.org/0000-0002-4347-8819>

## References

- [1] Behler J and Parrinello M 2007 *Phys. Rev. Lett.* **98** 146401
- [2] Shapeev A V 2016 *Multiscale Model. Simul.* **14** 1153
- [3] Thompson A P, Swiler L P, Trott C R, Foiles S M and Tucker G J 2015 *J. Comput. Phys.* **285** 316
- [4] Drautz R 2019 *Phys. Rev. B* **99** 014104
- [5] Artrith N and Urban A 2016 *Comput. Mater. Sci.* **114** 135
- [6] Bartók A P, Payne M C, Kondor R and Csányi G 2010 *Phys. Rev. Lett.* **104** 136403
- [7] Schütt K T, Sauceda H E, Kindermans P J, Tkatchenko A and Müller K R 2018 *J. Chem. Phys.* **148** 241722
- [8] Deringer V L, Bernstein N, Csányi G, Mahmoud C, Ceriotti M, Wilson M, Drabold D A and Elliott S R 2021 *Nature* **589** 59
- [9] The fitSNAP repository (available at: <https://github.com/FitSNAP/FitSNAP>)
- [10] Lilienfeld O A V, Ramakrishnan R, Rupp M and Knoll A 2015 *Int. J. Quantum Chem.* **115** 1084
- [11] Goryaeva A M, Maillet J-B and Marinica M-C 2019 *Comput. Mater. Sci.* **166** 200
- [12] Deringer V L, Bartók A P, Bernstein N, Wilkins D M, Ceriotti M and Csányi G 2021 *Chem. Rev.* **121** 10073
- [13] Li Z, Kermode J R and De Vita A 2015 *Phys. Rev. Lett.* **114** 096405
- [14] Vandermause J, Torrisi S B, Batzner S, Xie Y, Sun L, Kolpak A M and Kozinsky B 2020 *npj Comput. Mater.* **6** 20
- [15] Quiñonero-Candela J and Rasmussen C 2005 *J. Mach. Learn. Res.* **6** 1939
- [16] Snelson E and Ghahramani Z 2005 *Advances in Neural Information Processing Systems* **18**
- [17] Foster L, Waagen A, Aijaz N, Hurley M, Luis A, Rinsky J, Satyavolu C, Way M J, Gazis P and Srivastava A 2009 *J. Mach. Learn. Res.* **10** 857
- [18] The QUIP repository (available at: <https://github.com/libAtoms/QUIP>)
- [19] The extended XYZ format (available at: <https://github.com/libAtoms/extxyz>)
- [20] Mahoney M W and Drineas P 2009 *Proc. Natl Acad. Sci. USA* **106** 697
- [21] Rowe P, Deringer V L, Gasparotto P, Csányi G and Michaelides A 2020 *J. Chem. Phys.* **153** 034702
- [22] Csányi G 2022 Private communication
- [23] Bartók A P, Kermode J, Bernstein N and Csányi G 2018 *Phys. Rev. X* **8** 041048
- [24] Lejaeghere K et al 2016 *Science* **351** aad3000
- [25] 2022 Specifications of the raven cluster of the Max Planck Computing and Data Facility (MPCDF) (available at: <https://docs.mpcdf.mpg.de/doc/computing/raven-user-guide.html>)
- [26] Bartók A P, Kondor R and Csányi G 2013 *Phys. Rev. B* **87** 184115
- [27] Byggmästar J, Nordlund K and Djurabekova F 2021 *Phys. Rev. B* **104** 104101
- [28] Imbalzano G, Zhuang Y, Kapil V, Rossi K, Engel E A, Grasselli F and Ceriotti M 2021 *J. Chem. Phys.* **154** 074102