




Contents lists available at ScienceDirect

Computational Materials Science

journal homepage: www.elsevier.com/locate/commatsci

Full length article



Multi-objective optimization and quantum hybridization of equivariant deep learning interatomic potentials

G. Laskaris^{a,b}, D. Morozov^{a,c}, D. Tarpanov^a, A. Seth^d, J. Procelewska^e, G. Sai Gautam^d, A. Sagingalieva^a, R. Brasher^a, A. Melnikov^a^{*}^a Terra Quantum AG, Kornhausstrasse 25, 9000 St., Gallen, Switzerland^b LIACS, Leiden University, Netherlands^c Nanoscience Center and Department of Chemistry, University of Jyväskylä, P.O. Box 35, 40014 Jyväskylä, Finland^d Department of Materials Engineering, Indian Institute of Science, Bengaluru 560012, India^e Schaeffler Technologies AG & Co. KG, Herzogenaurach, Germany

ARTICLE INFO

Keywords:

Machine learning interatomic potentials

Multi-objective optimization

Equivariant neural networks

Quantum-classical hybrid models

Density functional theory

Copper-lithium interfaces

ABSTRACT

Allegro is a machine learning interatomic potential model designed to predict atomic properties in molecules using E(3) equivariant neural networks. When training this model, there tends to be a trade-off between accuracy and inference time. For this reason, we apply multi-objective hyperparameter optimization to both objectives. Additionally, we experiment with modified architectures by constructing variants of Allegro: one extended with additional classical layers and one incorporating quantum-classical hybrid layers. We evaluate all models on QM9, rMD17-aspirin, rMD17-benzene, and a self-generated dataset of copper-lithium structures. As results, both variants surpass Allegro in force prediction accuracy across multiple datasets. The classical variant consistently improves over the baseline, while the quantum-classical hybrid variant achieves the best overall force prediction accuracy on the Cu-Li dataset, where it was fully optimized, outperforming the classical variant by approximately 13%. Notably, the hybrid variant also achieves competitive results on the remaining datasets despite using hyperparameters transferred from Cu-Li without dataset-specific optimization, suggesting that quantum-classical hybridization is a promising direction for enhancing MLIP architectures.

1. Introduction

Machine learning interatomic potential (MLIP) models [1–3] are machine learning models trained to predict energies, stresses, forces of molecular systems. They map the atomic structure of molecules into their potential energies. An MLIP model, in comparison with density functional theory (DFT) [4,5], reduces computational cost and lowers the complexity of energy calculations, while approaching DFT accuracy. Some widely used MLIP models are SchNet [6], Gaussian Approximation Potentials (GAP) [7], Smooth Overlap of Atomic Positions - GAP (SOAP-GAP) [8], Atomic Cluster Expansion (ACE) [9,10] and NequiP model [11]. Here we conduct a methodical multi-objective optimization search of hyperparameters for an improved version of the aforementioned NequiP model which is the Allegro model [12]. Allegro yields significant improvements over NequiP [12], especially on large-scale systems. It allows for parallel computations by dropping the message passing propagation of the information and substituting it with a series of tensor products of learned equivariant representations. We optimize both the objectives of the accuracy and the inference

time. In addition, we perform a network architecture search for variants of the Allegro model, including quantum-classical hybrid models. We extend our hyperparameter optimization search on additional Allegro model variants, such as an extension of Allegro with additional multi-layer perceptron (MLP) layers along with quantum depth infused (QDI) layers [13–18]. There have been other areas where researchers have found success by inserting quantum layers in purely classical models, a process called, quantum-classical hybridization [13,19–22]. To ensure that we find the best fitting hyperparameters for all variant Allegro models, we use a state-of-the-art multi-objective optimization algorithm, our implementation of SAMO-COBRA [23] for the hyperparameter optimization. This algorithm runs multiple evaluations of our objective function and at the same time it searches and optimizes 6–9 hyperparameters depending on the type of Allegro model. We evaluate the performance of variant Allegro models on a plethora of datasets. One of them is an unpublished and self-generated dataset, and three widely used public datasets. The unpublished dataset includes inorganic compounds and specifically copper and lithium atoms. The public

^{*} Corresponding author.E-mail address: alexey@melnikov.info (A. Melnikov).<https://doi.org/10.1016/j.commatsci.2026.114742>

Received 18 February 2026; Received in revised form 21 April 2026; Accepted 25 April 2026

Available online 18 May 2026

0927-0256/© 2026 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

datasets are: (1) QM9 [24], which has 133885 different molecules, produced by all the stable combinations of up to 9 atoms of C, N, O and F, and (2) the revised-MD17 (rMD17) dataset, from which we use the Aspirin and Benzene subsets [25]. We train the models to predict molecular energies from QM9 and both energies and forces for Cu-Li, Aspirin and Benzene datasets.

2. Preliminaries

In this section, we outline key concepts of our work. We discuss the importance of equivariance and symmetries in graph neural networks (GNNs), introduce our implementation of SAMO-COBRA, and describe the QDI layer enhancements of the Allegro model.

2.1. Allegro model

In the context of molecular dynamics and the prediction of energies and forces, selecting an appropriate symmetry group G is vital. Equivariance arises naturally by the geometrical properties and symmetries of molecular systems [26], and thus comprises a mechanism that boosts the performance of MLIP [11,12]. The benefit of equivariance derives from the ability to reduce the space of possible weight matrices for a multilayer perceptron, in contrast with other MLIP models which require a large amount of configurations to achieve high-fidelity results [27,28]. Allegro model [12,29,30] is a Graph Neural Network (GNN) that utilizes internal symmetries to build an equivariance neural network [31–34] with a reduced number of trainable parameters. The demand for more accurate and sophisticated neural networks leads us to exploit different characteristics and properties. The notion of equivariance of a function to a group is characterized as a generalization of the invariance [31]. Regarding Allegro, the group over which the network is equivariant is $E(3)$. The group operations of $E(3)$ encompass rotations, translations in 3 dimensions, and reflections.

Definition 1 (Equivariance). Suppose the sets \mathcal{X}, \mathcal{Y} and a function $f : \mathcal{X} \rightarrow \mathcal{Y}$. Then for a group G and a transformation $g \in G$, if:

$$f(T_g x) = T'_g f(x), \quad \forall x \in \mathcal{X}, g \in G \quad (1)$$

Then f is equivariant to group G .

Note that T_g and T'_g are the representation operators of g in \mathcal{X} and \mathcal{Y} spaces, respectively. Also, these operators do not have to be equal but T has to be a linear representation of G [32]. In case of T'_g being the identity transformation then we fall to the special case of f being invariant to group G . For our case, the commutative diagram below, represents graphically the equivariance.

$$\begin{array}{ccc} \mathcal{X} & \xrightarrow{T_g} & \mathcal{X} \\ \downarrow f & & \downarrow f \\ \mathcal{Y} & \xrightarrow{T'_g} & \mathcal{Y} \end{array}$$

Definition 2 (k -Layered Multilayer Perceptron [31]). Let $\mathcal{A} \in \mathbb{R}^{n_{i-1} \times n_i}$, $i = \{1, \dots, k\}$ be all possible weights matrices, $\mathbf{b}_i \in \mathbb{R}^{n_i}$, $i = \{1, \dots, k-1\}$ be the bias terms (assuming no bias on the output layer) and $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ be a continuous activation function. Let G be a group with representation $\rho_i : G \rightarrow GL(n_i)$, $i = \{0, \dots, k\}$. Then the k -layered multilayer perceptron $f : \mathbb{R}^{n_0} \rightarrow \mathbb{R}^{n_k}$ is

$$f(x) = \mathcal{A}_k \sigma_{\mathbf{b}_{k-1}} \mathcal{A}_{k-1} \sigma_{\mathbf{b}_{k-2}} \dots \sigma_{\mathbf{b}_1} \mathcal{A}_1 x$$

Then f is equivariant to G with respect to ρ_i if

$$\mathcal{A}_i \rho_{i-1}(g) = \rho_i(g) \mathcal{A}_i, \quad \text{and} \quad (2)$$

$$\sigma_{\mathbf{b}_i} \rho_i(g) = \rho_i(g) \sigma_{\mathbf{b}_i}, \quad \forall g \in G \quad (3)$$

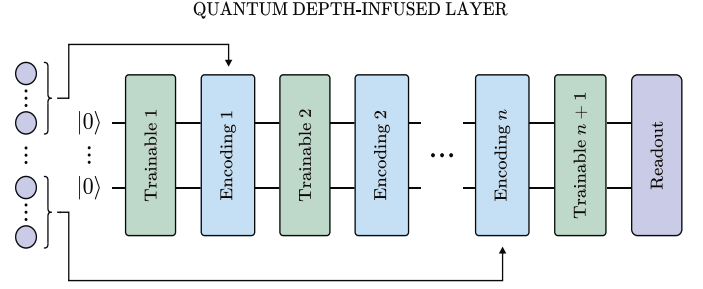


Fig. 1. Quantum depth-infused layer.

With $\sigma_{\mathbf{b}} : \mathbb{R}^{n_i} \rightarrow \mathbb{R}^{n_i}$, $\sigma_{\mathbf{b}}(x) = \sigma(x + \mathbf{b})$ being a point-wise nonlinearity.

Eqs. (2) and (3) are both derived from equivariance definition replacing the representation operator with the representation $\rho_i(g)$ of i th layer. Then by following the proof from [31] we can show that $f(\rho_0(g)x) = \rho_k(g)f(x)$. Eq. (3) is always satisfied by any point-wise nonlinearity while all components of \mathbf{b}_i are equal to the same real number. While this is true, it is evident how equivariance limits the possible choices for weights matrices, and thus reduces the complexity of our problem. As a consequence, neural networks that utilize the equivariance symmetry are capable on achieving high-fidelity results with significant less data [11,12].

2.2. Details of our implementation of SAMO-COBRA

We use our implementation of SAMO-COBRA as a multi-objective optimization algorithm that attempts to find the set of hyperparameters that minimize all the (expensive) objective functions we want to optimize. Particularly, our implementation of SAMO-COBRA succeeds on maximizing the hyper-volume in the solution space by minimizing the evaluations of the objective functions. Formally,

$$\begin{aligned} &\text{minimize } \mathbf{f}(\mathbf{x}) = \Omega \rightarrow \mathbb{R}^k, \mathbf{f}(\mathbf{x}) = \{f_1(\mathbf{x}), \dots, f_k(\mathbf{x})\} \\ &\text{subject to } g_i(\mathbf{x}) \leq 0, \forall i \in \{1, \dots, m\}, \mathbf{x} \in \Omega \subset \mathbb{R}^d \end{aligned} \quad (4)$$

where k is the number of objective functions we want to minimize, m is the number of constraints and d the dimensionality of the problem. In our case, we want to minimize two objective functions, namely mean absolute error (MAE) of forces with respect to the validation set and the inference time of the entire test set. We use our implementation of SAMO-COBRA to maximize the hypervolume (Hv) [23] of the region in the objective space dominated by a given set of solutions.

2.3. Quantum depth-infused layer

In the proposed ML model, a quantum layer, implemented as a specialized variational quantum circuit (VQC) known as the QDI layer [13], is leveraged to manage a large number of input features with a limited qubit count. The structure of this quantum layer is illustrated in Fig. 1. Unlike traditional approaches where each feature is assigned its own qubit, the QDI layer utilizes a data re-uploading strategy. This method sequentially encodes features onto a lattice structure. By incorporating entangling gates and variational quantum layers across encoding sections, the model can capture complex, non-linear patterns and correlations. This design enhances the model's representational power.

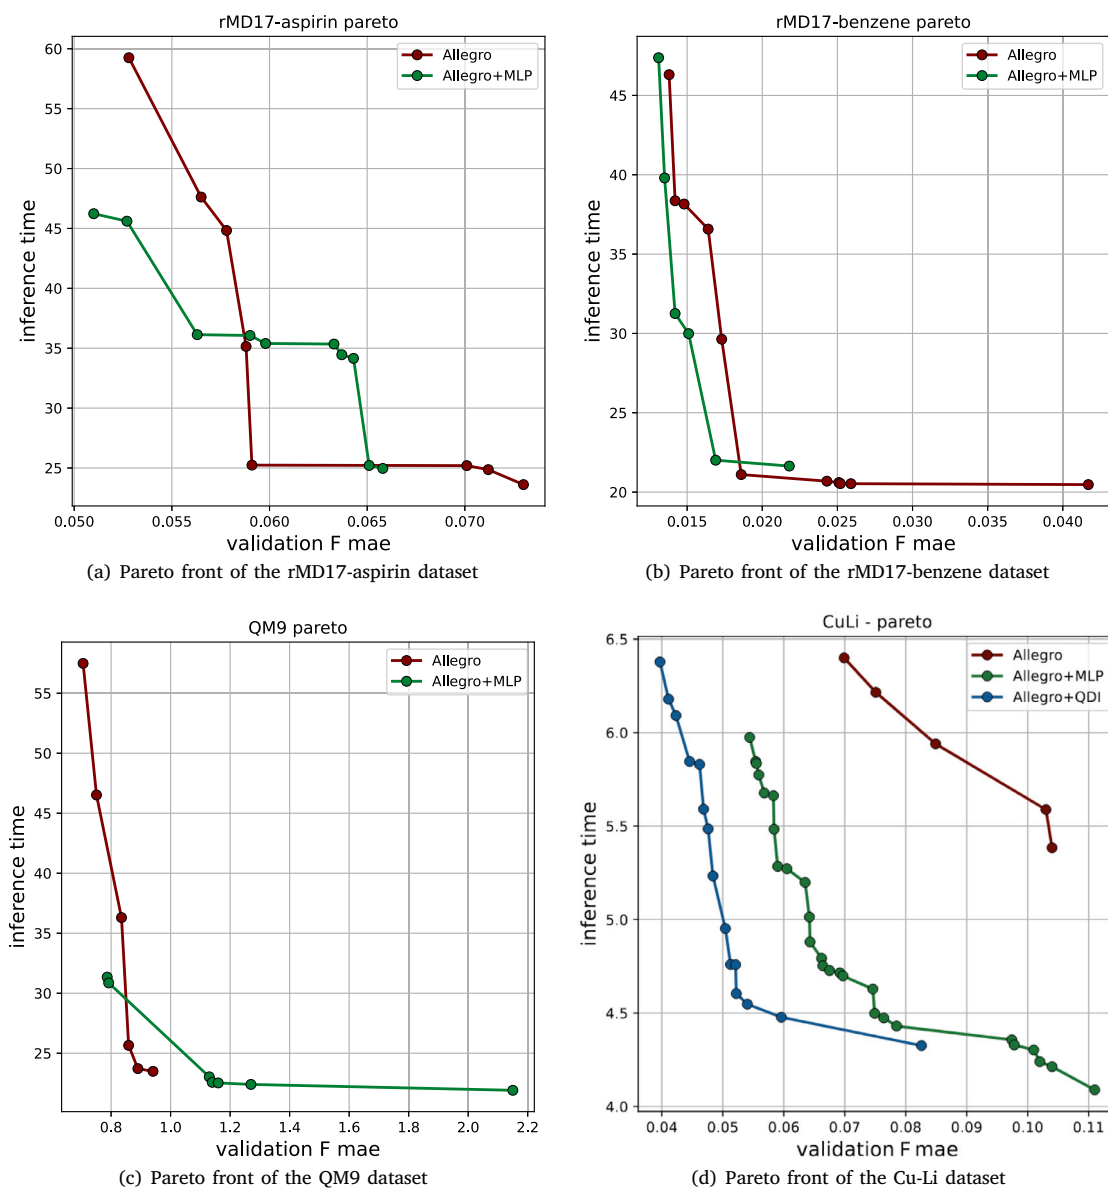


Fig. 2. Pareto fronts as they generated by our implementation of SAMO-COBRA during multi-objective optimization of the Allegro model and its variants, minimizing both the forces mean absolute errors on the validation set in ($\text{eV}/\text{\AA}$) and the inference time for each dataset tested. The x -axis shows the force mean absolute error (F MAE) on the validation set in $\text{eV}/\text{\AA}$, and the y -axis shows the total inference time on the entire test set in seconds.

3. Implementation

For the implementation and the experiments of this work, we generated our own Cu-Li dataset from DFT calculations, we preprocessed all datasets before running any training, built all the necessary variant Allegro models, namely Allegro+MLP and Allegro+QDI and used a single T4 GPU with 16 GB of memory to run the multi-objective optimization using our implementation of SAMO-COBRA as our optimization algorithm.

3.1. Data preprocessing

The choice of the datasets serves multiple purposes. Training the models on QM9 allows us to explore the potential of these predictive models on a diverse set of organic molecules. We pre-processed QM9 by excluding the 3054 molecules which failed the geometry consistency check from the total 133885 molecules [24]. Additionally, from all

molecules' features, we isolated the energy values used for the training. For more focused experiments we used rMD17 and specifically the aspirin and benzene subsets. For consistency with prior work, we converted all the energies from kcal/mol to eV/atom , since we compare our results with the literature where energies are usually reported in eV or meV . Finally, to evaluate Allegro model and its variants on inorganic compounds, we generated a dataset consisting of copper and lithium atoms. This dataset consists of multiple XYZ -formatted slabs which form a surface of copper atoms and their respective energy.

3.2. Dataset generation

The DFT-based training dataset was constructed using a diverse set of 11,635 structures, enabling the trained MLIP to accurately model complex materials behavior in different systems. We first generated an initial set of 13 symmetrically distinct slabs using *pymatgen* [35] for Li and Cu each up to the {333} Miller indices. The thickness of the

Table 1
Best parameters from Allegro HPO.

Allegro	r_{max}	lr	Batch	P	Layers	Edge MLP
rMD17 aspirin	7.3	5.36e-03	3	6.9	4	32
rMD17 benzene	12.0	1e-03	1	5.0	3	128
QM9	12.0	1e-03	5	10.0	3	32
CuLi	12.0	2e-03	2	5.0	2	128

Table 2
Best parameters from Allegro+MLP HPO.

Allegro+MLP	r_{max}	lr	Batch	P	Layers	Edge MLP	Final MLP
rMD17 aspirin	7.4	2.93e-03	2	6.7	3	16	[1024,1024]
rMD17 benzene	12.0	1e-03	1	5.0	3	128	[32,4]
QM9	6.2	1e-03	2	5.0	2	64	[32,16]
CuLi	5.0	9.35e-04	1	10.3	3	16	[4,16]

Li and Cu slabs was fixed as 20 Å and 30 Å by convergence of their surface energies respectively with a vacuum spacing of 15 Å on both sides to prevent periodic boundary interactions. To distinguish the bulk behavior of the slab from its surface, we only allow the atoms in the first and last two layers of each slab to relax during the DFT calculations, while immobilizing the rest of the atoms. These relaxed configurations were then used to create a wide variety of slab and interface-based structures. Firstly, we generate structures with surface vacancies and adatoms to capture the effect of surface modifications on these slabs. 34 structures with surface vacancies were created by removing the furthest possible surface atoms having the same coordination number in each slab. To capture defect concentration, we further generated surface vacancies in $2 \times 2 \times 1$ supercells of the slabs leading to an additional 34 structures. 537 structures with adatoms were generated by placing Li and Cu adatoms at distinct sites (hollow, top and bottom) at 2 Å on top of the $2 \times 2 \times 1$ supercells of each slab and allowing the structures to relax in order to capture the adsorption behavior on the surfaces. Lastly, we generated a set of interfaces and melt-quench based amorphous interfacial structures to enable the potential to accurately model high-temperature environments with complex interactions between Li and Cu atoms. We used *pymatgen* to create coherent interfaces between the Li and Cu slabs following the approach suggested by Zur and McGill [36]. We started by initially transforming Li and Cu slabs into their respective supercells such that their surface areas are similar within a defined tolerance of 9%. For each pair of Li and Cu slabs, we found the smallest possible combination of such supercells and positioned them with a gap of 3, thus creating geometrically viable interfaces. Due to the enormous number of interfaces that can be generated via this process, we constrained our dataset to having interfaces generated by the Li and Cu slabs sharing the same Miller indices. A melt-quench approach via ab initio molecular dynamics (AIMD) simulations was subsequently used to create amorphized interfacial structures.

For the melt-quench simulations, the interfaces were heated to 600K, which is above the melting point of Li (~ 453 K) but well below that of bulk Cu (~ 1358 K). At this temperature, the Li component melts while the Cu surface layers undergo significant disorder due to interfacial mixing with molten Li, producing amorphous interfacial structures representative of realistic Li-Cu interfaces encountered, for example, in lithium-ion battery anodes. The system was then cooled from 600 K to 10 K at a rate of 250K/ps. The NVT ensemble with the Nose-Hoover thermostat [37–39] was used for all the AIMD simulations with a time step of 2 fs and the energy cutoff fixed as 520 eV. This interfacial data forms the bulk of our dataset and consists of 11,000 distinct structures.

3.3. DFT calculations

We utilized the Vienna ab initio simulation package (VASP) [40,41] with the projector augmented wave [42,43] potentials to generate the entire dataset. The restored-regularized strongly constrained and appropriately normed (R²SCAN) [44] exchange correlation functional was used to model the electronic exchange and correlation, with the kinetic energy cutoff fixed as 520 eV for the plane wave basis. We sampled the Brillouin zone using Γ centered Monkhorst-Pack k -point meshes having a minimum density of $32/\text{Å}$. We allowed for the ionic positions to be relaxed for all configurations without symmetry constraints until the total energies and atomic forces for all structures converged within 10^{-5} eV and $|0.03|$ eV/Å, respectively.

3.4. Allegro enhancements

The first modification made to the existing Allegro structure was the insertion of an additional MLP layer (see Allegro architecture figure in [12]). This MLP consists of two hidden layers, the dimensionality of which (along with the input number of features) consists of three additional hyperparameters for tuning. Moreover, the weights initialization and nonlinearity methods constitute two hyperparameters that we included to the configuration of the model but did not optimize. These MLP layers add some additional depth to the model, which increases the complexity and ultimately leads to better accuracy. Allegro+MLP serves as a direct classical counterpart to Allegro+QDI allowing for a proper comparison between the two models. This comparison is valid since both Allegro+MLP and Allegro+QDI share the same overall structure. The only difference is that the MLP, including the two hidden layers, is replaced by a QDI layer followed by a fully connected layer, before it passes to the final output MLP of the original Allegro model.

3.5. Hyperparameter tuning

For the hyperparameter optimization of all Allegro, Allegro+MLP and Allegro+QDI models we used our implementation of SAMO-COBRA which evaluated the objective function for 100 generations and 12 off-springs, implying a total of 1200 evaluations for each variant and each dataset. For transparency to computational costs, we report the mean wall-clock time per expensive evaluation on a single T4 GPU at approximately 32.4 min for QM9, 9.6 min for rMD17-aspirin, 6.0 min for rMD17-benzene and 4.8 min for Cu-Li. Multiplying with the number of expensive evaluations of the optimization process, gives an estimate of the total optimization cost per dataset. We optimized Allegro on 6 different hyperparameters, including the cutoff radius, the polynomial cutoff, batch size, number of layers, dimensionality of the final MLP layer and learning rate. Especially for the Allegro+MLP model, we optimized also the dimensionality of the two hidden layers and the number of input features. For Allegro+QDI, the additional optimized hyperparameters where the number of input features to the quantum circuit, the dimensionality of the additional MLP and the depth of the QDI. The models were optimized to the minimize mean absolute error on forces (mean absolute error on energies for QM9) for the validation set, and at the same time the minimization of the inference time of the models on a test set. To visualize and evaluate the optimization, we produced the Pareto fronts for all given datasets Fig. 2. Unfortunately, we were not able to apply the full Pareto-front optimization on Allegro+QDI for the rMD17-aspirin, rMD17-benzene, and QM9 datasets. Consequently, we produced the complete Pareto front for Allegro+QDI only on Cu-Li. For the remaining datasets, we transferred the best QDI-specific hyperparameters (QDI depth and QDI input features) found on Cu-Li, while using the dataset-specific Allegro hyperparameters optimized for each respective dataset. We note that this transfer likely represents a lower bound on Allegro+QDI performance, since dataset-specific optimization of the QDI hyperparameters could yield further improvements.

Table 3
Best parameters from Allegro+QDI HPO.

Allegro+QDI	r_{\max}	lr	Batch	P	Layers	Edge MLP	Final MLP	QDI depth	QDI input features
rMD17 aspirin	9.5	1e-03	1	6.0	4	16	4	4	21
rMD17 benzene	12.0	1e-03	1	5.0	3	32	16	1	4
QM9	12.0	1e-03	5	10.0	3	32	128	3	4
CuLi	5.0	4.02e-03	5	12.0	3	4	128	3	4

Table 4
 $\text{Val}_{f_{\max}}$ values for best runs in (meV/Å).

$\text{Val}_{f_{\max}}$	Allegro	Allegro [12]	Allegro+MLP	Allegro+QDI
Aspirin	10.2	9.8	9.0	9.9
Benzene	0.43	0.37	0.28	0.30
CuLi	45.9	–	43.3	37.8

Table 5
 $\text{Val}_{e_{\max}}$ values for best runs in (meV/atom) for up to 300 epochs on QM9, 100 for CuLi and 3000 for aspirin and benzene.

$\text{Val}_{e_{\max}}$	Allegro	Allegro [12]	Allegro+MLP	Allegro+QDI
Aspirin	3.6	2.3	2.6	3.2
Benzene	0.21	0.20	0.21	0.19
CuLi	809.0	–	698.9	495.7
QM9	10.3	13.4	21.4	23.3

4. Results

By running the multi-objective optimization on all Allegro variants and all datasets we plot the respective Pareto fronts in Fig. 2. For all datasets, there is a trade-off between accuracy and inference time. For rMD17-aspirin and rMD17-benzene we notice that Allegro+MLP can achieve better accuracy but pure Allegro validates a test set faster, while for QM9 it is the opposite. Finally, CuLi dataset in Fig. 2 includes the Allegro+QDI Pareto front, which we manually calculated due to technical reasons with the inference time and the optimization library that we use. In principle, it is possible to modify our approach for the QDI variant to make it generally efficient. Nevertheless, for CuLi, Allegro+QDI has the overall best accuracy on predicting the forces on the validation set and slightly worse inference time in comparison with Allegro+MLP. In Tables 1, 2 and 3 we include all the best hyperparameters found by our implementation of SAMO-COBRA on vanilla Allegro, Allegro+MLP and Allegro+QDI respectively. We notice, large differences in specific hyperparameters (such as the polynomial envelope P and the dimensionality of the edge MLP), between the different variants, which is an indication of a non-significant hyperparameter. Due to the aforementioned technical issue, we optimize Allegro+QDI only on the forces MAE. Specifically for the QM9 dataset, instead of optimizing the Allegro+QDI model we use the optimized hyperparameters on QM9 from Allegro and for the QDI-exclusive hyperparameters, we use the optimized Allegro+QDI parameters from CuLi. In Tables 4 and 5 we report the results from training the models on all datasets,

using the optimized hyperparameters. In the following supplementary material, we include all the graphs from the training of these runs. We compare all our results with the results reported in the original Allegro paper [12]. Since we were unable to replicate the exact configuration file as the one used in [12], we used a common configuration file for all models and datasets and only changed the optimized hyperparameters in each case. In the tables, we denote with bold the most accurate models, which in all cases (for forces) are the variant Allegro+MLP or Allegro+QDI. The Allegro+QDI variant dominates on the CuLi dataset, offering almost 13% increase in performance over Allegro+MLP. Table 6 places our results in the context of other state-of-the-art MLIP models. On rMD17-aspirin, our Allegro+MLP variant (9.0 meV/Å) achieves competitive force MAE, outperforming the original Allegro and BOTNet, though MACE achieves a lower error (6.6 meV/Å). On rMD17-benzene, Allegro+MLP (0.28 meV/Å) is competitive with NequIP and MACE. We emphasize that the primary contribution of this work is not to claim state-of-the-art accuracy on these benchmarks, but rather to demonstrate the value of systematic multi-objective hyperparameter optimization and to explore the potential of quantum-classical hybridization within a specific MLIP framework. Direct comparisons should also account for differences in training configurations, data splits, and computational resources. Regarding energies, we notice a more unstable behavior during training which is justified in such accurate calculations (in the range of a few meV and less than 1 meV for benzene). Since the optimization focuses only on forces and inference time, the energy results do not highlight any model to be strictly superior.

5. Discussion

To summarize, we use our implementation of SAMO-COBRA to optimize hyperparameters on an Allegro model along with the classical and hybrid variants we built. The optimization targeted two objectives: minimizing the mean absolute error on validation set forces and minimizing inference time. The results indicate that the classical variant Allegro+MLP consistently surpasses vanilla Allegro in force prediction accuracy across all tested datasets. The quantum-classical hybrid Allegro+QDI achieves the best force prediction accuracy on the CuLi dataset, where it was fully optimized, offering approximately 13% improvement over Allegro+MLP. Notably, even with hyperparameters transferred from Cu-Li, Allegro+QDI shows competitive performance on the remaining datasets and achieves the best energy MAE on rMD17-benzene, suggesting that the quantum hybridization benefits are not

Table 6
Comparison with state-of-the-art MLIP models on rMD17 and QM9 benchmarks. Values for NequIP, MACE, and BOTNet are taken from Table 1 of Ref. [45], trained on 950 configurations with 50 for validation.

Model	rMD17-aspirin		rMD17-benzene		QM9
	F_{MAE} (meV/Å)	E_{MAE} (meV/atom)	F_{MAE} (meV/Å)	E_{MAE} (meV/atom)	E_{MAE} (meV/atom)
NequIP [11]	8.2	2.3	0.3	0.04	–
MACE [45]	6.6	2.2	0.3	0.4	–
BOTNet [46]	8.5	2.3	0.3	0.03	–
Allegro [12]	7.3	2.3	0.2	0.3	13.4
Allegro (ours)	10.2	3.6	0.43	0.21	10.3
Allegro+MLP (ours)	9.0	2.6	0.28	0.21	21.4
Allegro+QDI (ours) ¹	9.9	3.2	0.30	0.19	23.3

¹ Allegro+QDI results on rMD17 and QM9 use transferred hyperparameters from Cu-Li (see text).

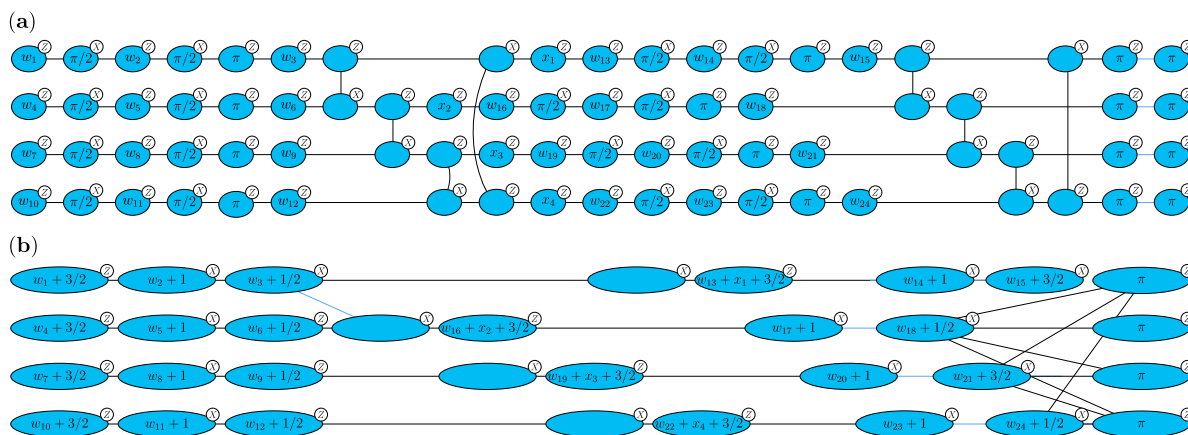


Fig. 3. (a) Circuit's ZX graph with original parameters. (b) Circuit's ZX graph after simplification.

limited to a single dataset. For future work, one could explore the capabilities of such state-of-the-art MLIP models (e.g., Allegro), on more complex molecules that are impossible to solve analytically, where such highly accurate, approximation methods can prove significant utility. Based on our work here, we expect the hybridization of such MLIPs to offer additional accuracy benefits. Thus, exploring different ansatzes or involving more classical/quantum layers in the architecture of the model can improve them significantly.

CRedit authorship contribution statement

G. Laskaris: Writing – review & editing, Writing – original draft, Investigation, Data curation. **D. Morozov:** Software, Investigation, Formal analysis. **D. Tarpanov:** Software, Investigation. **A. Seth:** Writing – review & editing, Validation, Investigation. **J. Procelewska:** Writing – review & editing, Writing – original draft, Data curation, Conceptualization. **G. Sai Gautam:** Writing – review & editing, Writing – original draft, Methodology, Investigation. **A. Sagingalieva:** Writing – review & editing, Writing – original draft, Software, Conceptualization. **R. Brasher:** Writing – review & editing, Software, Investigation, Formal analysis. **A. Melnikov:** Writing – review & editing, Supervision, Project administration.

Declaration of competing interest

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

Supplementary material A. Quantum circuits analysis

In this section, we analyze the quantum circuit used in this paper. There are three metrics for evaluation:

- Redundancy analysis with ZX calculus
- Trainability analysis with Fisher information
- Expressivity analysis with Fourier series

A.1. Redundancy analysis with ZX calculus

ZX calculus is a graphical language used to simplify quantum circuits using specific notation, where the circuit is transformed into a sequence of nodes, called “spiders”, and edges to connect them. After this ZX graph can be simplified [47] using language’s set of rules, thus taking out redundant elements [48]. After transformation, we can compare the initial version with the reduced one, which represents a more optimal configuration. Using the number of parameters in the first and second circuits, the redundancy can be evaluated as the margin

of initial parameters left after simplification. The more parameters are left unchanged, the better architecture the circuit has in the first place. A architecture deemed unable to be simplified in this manner is classified as ZX-irreducible. The main changes in the circuit shown in Fig. 3(a–b) consist of the rearrangement of certain weights after their simplification. During the optimization process, 24 out of 24 parameters, or 100%, were preserved, showing the circuit’s high quality and lack of redundant elements. Using ZX calculus, it can be clearly said that a circuit does not need major structural changes and has no redundant elements in it. However, other metrics should be used to provide a full analysis of its capabilities.

A.2. Trainability analysis with Fisher information

A supervised machine learning task can be framed as constructing a hypothesis model $h_\theta(\hat{x})$ based on a labeled dataset $(x, y) \in X \times Y$ to approximate the natural data distribution $f(x)$. Given a subset S of labeled samples from this distribution, we aim to optimize our model to accurately approximate $f(\hat{x})$. This involves maximizing the probability of the model, with parameters θ , predicting the correct label y for each data point x . The conditional probability required for this task can be expressed as $P(y|x, \theta)$. By assuming a uniform distribution over X , we leverage the joint probability $P(y, x|\theta)$ to enhance accuracy, with this probability distribution calculable for any parameter value θ given a data point x_i . This joint probability can be represented as an N -dimensional manifold, where N is the number of trainable parameters, $N = |\theta|$. The Fisher information matrix (FIM) $F(\theta)$ [49,50] is a metric over this manifold.

$$F(\theta) = \mathbb{E}_{\{x_i, y_i\}} [\nabla_\theta \log(P) \nabla_\theta \log(P)^T] \quad (5)$$

To proceed, we diagonalize the metric to obtain a locally Euclidean tangential basis, where the diagonal elements represent the squared gradients of our joint probability within this basis. These elements correspond to the eigenvalues of the Fisher information matrix. This process is crucial for identifying and mitigating the barren plateau problem, which manifests as vanishing gradients in quantum neural networks with a high number of qubits. As demonstrated in Ref. [51], the expectation values of gradients approach zero, and their variance diminishes exponentially as the number of qubits increases. This effect arises when gradients predominantly converge near zero, leading to non-participation of many parameters in the training process. Consequently, analyzing the eigenvalue spectrum of Fisher matrices across multiple instances of θ provides insights into the trainability and resilience of QNNs against barren plateaus. A neural network with higher trainability would exhibit reduced eigenvalue degeneracy. The Fisher information matrix can be computed for the specific hyperparameters of the circuit. Following the approach outlined in Ref. [49], we generate

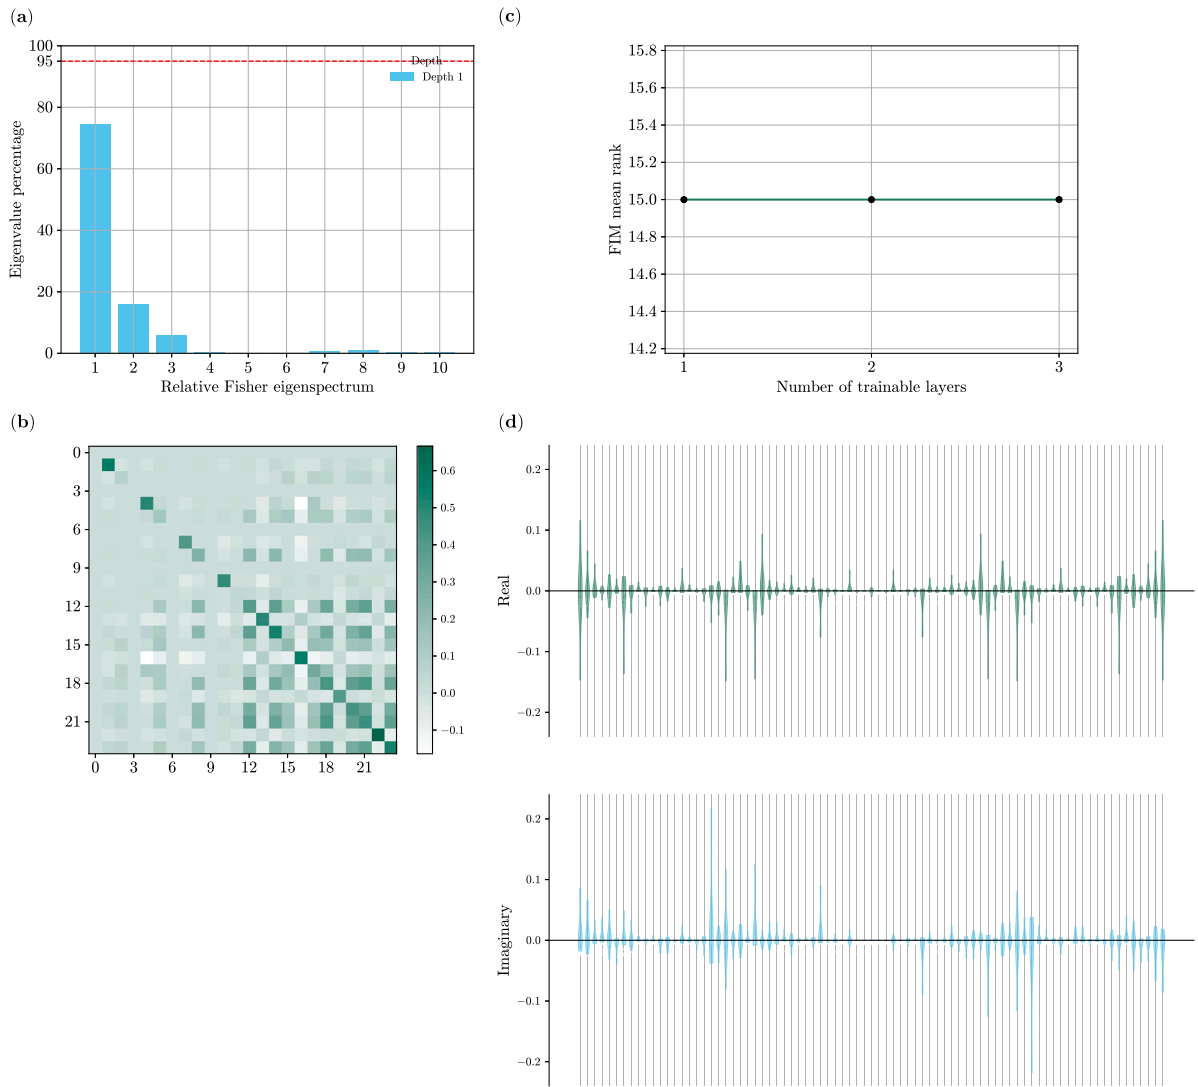


Fig. 4. (a) The histogram of the Fisher eigenspectrum. There are around 78% of parameters with near-zero eigenvalues, but this margin does not go higher than 95% trainability threshold. Other parameters have great impact, contributing to circuit's trainability. (b) Average Fisher Information Matrix. The diagonal elements show parameters having mostly high gradient values, showing their impact on the circuit. Meanwhile non-diagonal ones are almost non-existent, indicating parameters' independence. Also there is no evident single-parameter dominance, showing high trainability. (c) Average FIM rank. There are three layers presented, although only one is used in the original circuit. This clearly shows how rank goes on the plateau, which means the model is in the overparameterization mode. Thus, no extra layers are needed as the circuit already has a sufficient number of parameters. (d) Real and imaginary part of Fourier coefficients.

a Gaussian-distributed dataset $x \sim \mathcal{N}(\mu = 0, \sigma^2 = 1)$. The joint probability is then determined by evaluating the overlap between the prepared quantum state and the state produced by our quantum layer.

$$P(y, x|\theta) = \langle y|\psi(\theta, x)\rangle, \quad (6)$$

where y is the output state. By averaging over all x and y , the Fisher information is calculated for any given θ . The metric's result consists of several images. Fig. 4(a) shows distribution with several parameters having eigenvalues much greater than zero. This indicates that many of them are impactful, giving good trainability. In Fig. 4(b), we can see a clear diagonal where every element is active and contributes to the circuit's performance. Lack of non-diagonal elements also shows that there is no pair of parameters that are entangled to the extent where we can replace them with a singular one.

As discussed in Ref. [53], certain quantum neural networks (QNNs) may exhibit reduced parameter efficiency due to over-parametrization. This effect is observed when, beyond a certain point, adding more parameters no longer increases the rank of the FIM, indicating that the

circuit has reached a saturation point. At this stage, further parameter addition does not enhance expressivity and may lead to the risk of over-parametrization. Choosing the right number of parameters, enough to utilize the full capacity of the circuit but not too much, is the important issue that is resolved with the average FIM rank Fig. 4(c).

A.3. Expressivity analysis with fourier series

It has been demonstrated that the output of a parameterized quantum circuit can be represented as a truncated Fourier series. For a feature vector of length N , this Fourier series is expressed as a function of the feature vector x and the trainable parameters θ :

$$f_{\theta}(x) = \sum_{\omega_1 \in \Omega_1} \dots \sum_{\omega_N \in \Omega_N} c_{\omega_1, \dots, \omega_N}(\theta) e^{-i\omega x} \quad (7)$$

where $\omega \in \{d_{-i}, \dots, 0, \dots, d_i\}$. In other words, the number of terms in the Fourier series is one plus twice the number of times the input is embedded in the circuit, denoted by d . In this analysis, we assess the expressivity of the function $f_{\theta}(x)$ by sampling over a uniform

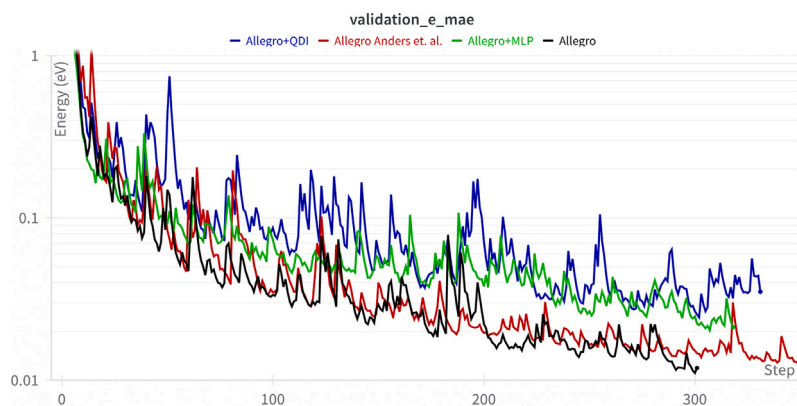


Fig. 5. Training curves of all Allegro variants on QM9 dataset using the best hyperparameters, as achieved from our hyperparameter optimization. Among the variants, we include the Allegro Anders et al. training curve, which corresponds to QM9 Allegro hyperparameters as reported in [12]. To smooth the resulting curves we use the time-weighted exponential moving average (EMA) [52] smoothing technique with 0.75 smoothing factor.

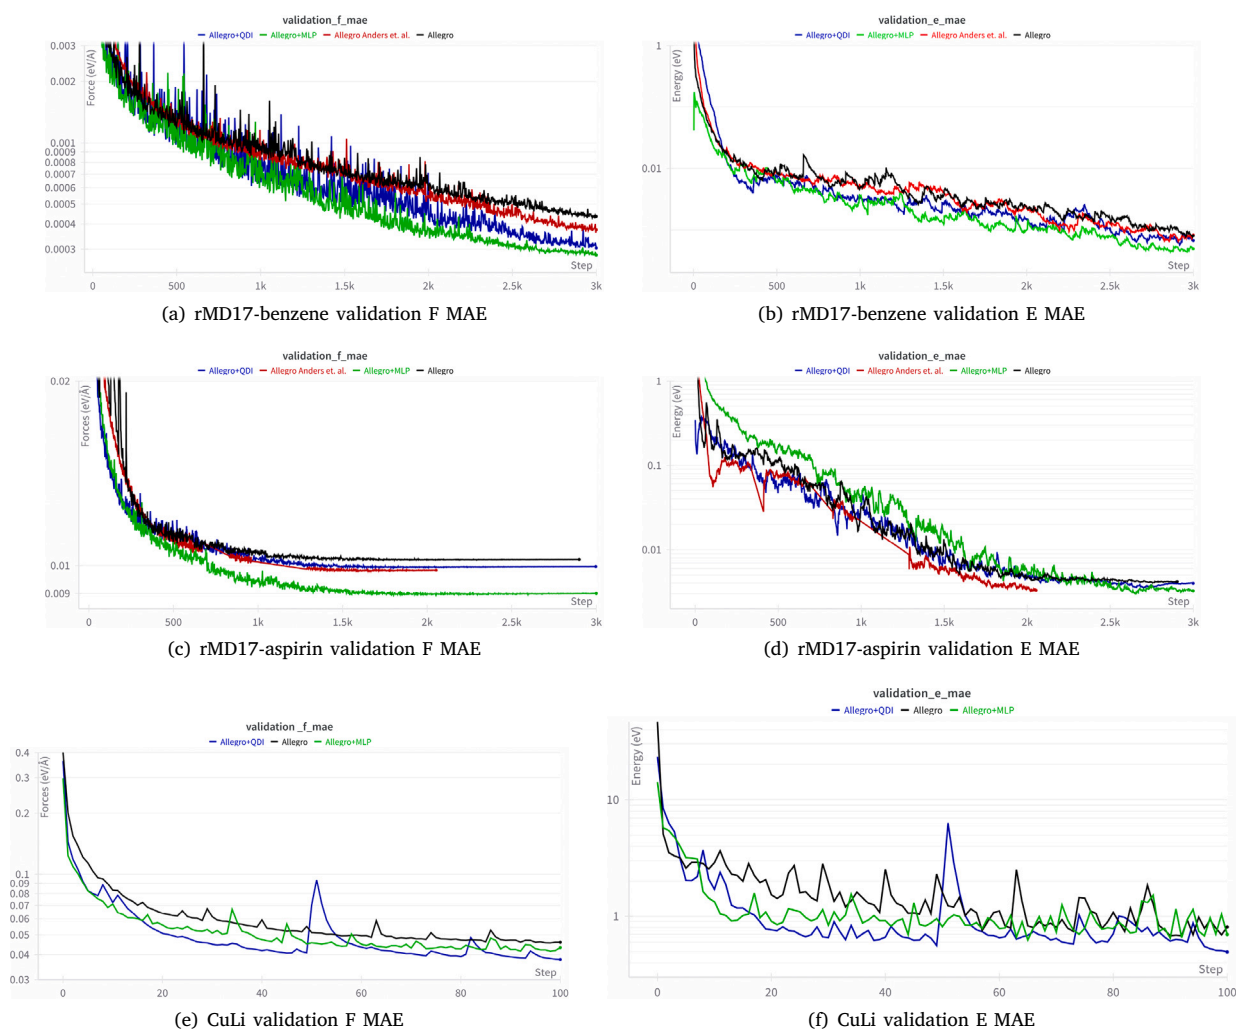


Fig. 6. Training curves for all Allegro variants including Allegro with the parameters reported by Anders et al. [12]. We report mean absolute errors both forces and energies for all datasets up to 3000 epochs (except from QM9 and CuLi which is up to 300 and 100 epochs respectively). Moreover, for the energy sub-figures (b) and (d) we use the time weighted EMA [52] smoothing technique with smoother factors 0.9 and 0.8 respectively. All graphs are plotted with semi-logarithmic scaling in y-axis.




Fig. 7. An example generated (a) Li (110) and (b) Cu (110) slabs with a thickness of 20 Å and 30 Å and vacuum spacing of 15 Å. Interface generated between the (c) Li (110) and Cu (110) slabs and (d) Li (111) and Cu (111) slabs with the same vacuum spacing of 15 Å.

distribution of random values for each θ_i within the interval $[0, 2\pi]$ and by selecting equidistant x values with a sampling frequency of d . This analysis centers on two primary considerations: the number of accessible terms in the Fourier series (the degree) and the accessibility of the coefficients for each term (the coefficients' expressivity). The expressivity is indicated by the number of non-zero coefficients in the Fourier spectrum—the higher the number, the greater the expressivity. Notably, this circuit demonstrates a high degree of expressivity, with 127 non-zero coefficients out of 161 possible terms (around 79%) as shown in Fig. 4(d).

Supplementary material B. Training curves

See Figs. 5–7.

Data availability

The datasets used in this study are publicly available. The rMD17 dataset is available as described in [25], and the QM9 dataset in [54]. The Cu-Li dataset was generated in collaboration with an industrial partner and is subject to data sharing restrictions. It is available upon reasonable request to the corresponding author for academic research purposes, subject to data sharing agreement.

The machine learning models and code used in this work include NequIP [11], Allegro [12], Allegro+MLP [55], and SAMO-COBRA [56], all of which are accessible via their respective GitHub repositories.

References

- [1] V.L. Deringer, M.A. Caro, G. Csányi, Machine learning interatomic potentials as emerging tools for materials science, *Adv. Mater.* 31 (46) (2019) 1902765, <http://dx.doi.org/10.1002/adma.201902765>.
- [2] D.M. Anstine, O. Isayev, Machine learning interatomic potentials and long-range physics, *J. Phys. Chem. A* 127 (11) (2023) 2417–2431, <http://dx.doi.org/10.1021/acs.jpca.2c06778>.
- [3] T. Mueller, A. Hernandez, C. Wang, Machine learning for interatomic potential models, *J. Chem. Phys.* 152 (5) (2020) 050902, <http://dx.doi.org/10.1063/1.5126336>.
- [4] N. Argaman, G. Makov, Density functional theory: An introduction, *Am. J. Phys.* 68 (1) (2000) 69–79, <http://dx.doi.org/10.1119/1.19375>.
- [5] M. Orto, D.A. Pantazis, F. Neese, Density functional theory, *Photosynth. Res.* 102 (2009) 443–453, <http://dx.doi.org/10.1007/s11120-009-9404-8>.
- [6] K.T. Schütt, H.E. Sauceda, P.-J. Kindermans, A. Tkatchenko, K.-R. Müller, SchNet—a deep learning architecture for molecules and materials, *J. Chem. Phys.* 148 (24) (2018) 241722, <http://dx.doi.org/10.1063/1.5019779>.
- [7] A.P. Bartók, M.C. Payne, R. Kondor, G. Csányi, Gaussian approximation potentials: The accuracy of quantum mechanics, without the electrons, *Phys. Rev. Lett.* 104 (13) (2010) 136403, <http://dx.doi.org/10.1103/PhysRevLett.104.136403>.
- [8] V.L. Deringer, A.P. Bartók, N. Bernstein, D.M. Wilkins, M. Ceriotti, G. Csányi, Gaussian process regression for materials and molecules, *Chem. Rev.* 121 (16) (2021) 10073–10141, <http://dx.doi.org/10.1021/acs.chemrev.1c00022>.
- [9] R. Drautz, Atomic cluster expansion for accurate and transferable interatomic potentials, *Phys. Rev. B* 99 (1) (2019) 014104, <http://dx.doi.org/10.1103/PhysRevB.99.014104>.
- [10] Y. Lysogorskiy, C. van der Oord, A. Bochkarev, S. Menon, M. Rinaldi, T. Hamerschmidt, M. Mrovec, A. Thompson, G. Csányi, C. Ortner, R. Drautz, Performant implementation of the atomic cluster expansion (PACE) and application to copper and silicon, *Npj Comput. Mater.* 7 (1) (2021) 97, <http://dx.doi.org/10.1038/s41524-021-00559-9>.
- [11] S. Batzner, A. Musaelian, L. Sun, M. Geiger, J.P. Mailoa, M. Kornbluth, N. Molinari, T.E. Smidt, B. Kozinsky, E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials, *Nat. Commun.* 13 (1) (2022) 2453, <http://dx.doi.org/10.1038/s41467-022-29939-5>.
- [12] A. Musaelian, S. Batzner, A. Johansson, L. Sun, C.J. Owen, M. Kornbluth, B. Kozinsky, Learning local equivariant representations for large-scale atomistic dynamics, *Nat. Commun.* 14 (1) (2023) 579, <http://dx.doi.org/10.1038/s41467-023-36329-y>.
- [13] A. Sagingalieva, M. Kordzanganeh, N. Kenbayev, D. Kosichkina, T. Tomashuk, A. Melnikov, Hybrid quantum neural network for drug response prediction, *Cancers* 15 (10) (2023) 2705, <http://dx.doi.org/10.3390/cancers15102705>.
- [14] A. Sagingalieva, S. Komorniyk, A. Senokosov, A. Joshi, C. Mansell, O. Tsurkan, K. Pinto, M. Pflitsch, A. Melnikov, Photovoltaic power forecasting using quantum machine learning, *Sol. Energy* 302 (2025) 114016, <http://dx.doi.org/10.1016/j.solener.2025.114016>.
- [15] L. Lusnig, A. Sagingalieva, M. Surmach, T. Protasevich, O. Michiu, J. McLoughlin, C. Mansell, G. de'Petris, D. Bonazza, F. Zanconati, A. Melnikov, F. Cavalli, Hybrid quantum image classification and federated learning for hepatic steatosis diagnosis, *Diagnostics* 14 (5) (2024) 558, <http://dx.doi.org/10.3390/diagnostics14050558>.

- [16] A. Sedykh, M. Podapaka, A. Sagingalieva, K. Pinto, M. Pflitsch, A. Melnikov, Hybrid quantum physics-informed neural networks for simulating computational fluid dynamics in complex shapes, *Mach. Learn.: Sci. Technol.* 5 (2) (2024) 025045, <http://dx.doi.org/10.1088/2632-2153/ad43b2>.
- [17] M. Anoshin, A. Sagingalieva, C. Mansell, D. Zhiganov, V. Shete, M. Pflitsch, A. Melnikov, Hybrid quantum cycle generative adversarial network for small molecule generation, *IEEE Trans. on Quantum Eng.* 5 (2024) 2500514, <http://dx.doi.org/10.1109/TQE.2024.3414264>.
- [18] N. Lee, M. Shin, A. Sagingalieva, A. Senokosov, M. Anoshin, A.J. Tripathi, K. Pinto, A. Melnikov, Predictive control of blast furnace temperature in steelmaking with hybrid depth-infused quantum neural networks, 2025, arXiv preprint [arXiv:2504.12389](https://arxiv.org/abs/2504.12389).
- [19] A. Sagingalieva, L. Lusnig, F. Cavalli, A. Melnikov, Hybrid quantum neural networks for computer-aided sex diagnosis in forensic and physical anthropology, *Informatics Med. Unlocked* 58 (2025) 101682, <http://dx.doi.org/10.1016/j.imu.2025.101682>.
- [20] O. Tsurkan, A. Konstantinova, A. Sedykh, A. Senokosov, D. Tarpanov, M. Anoshin, A. Sagingalieva, A. Melnikov, Hybrid quantum recurrent neural network for remaining useful life prediction, 2025, arXiv preprint [arXiv:2504.20823](https://arxiv.org/abs/2504.20823).
- [21] A. Kurkin, J. Hegemann, M. Kordzanganeh, A. Melnikov, Forecasting steam mass flow in power plants using the parallel hybrid network, *Eng. Appl. Artif. Intell.* 160 (2025) 111912, <http://dx.doi.org/10.1016/j.engappai.2025.111912>.
- [22] V. Lopatkin, A. Sagingalieva, L. Lusnig, T. Protasevich, B. Behnke, A. Melnikov, Quantum hybrid feature selector, *EPJ Quantum Technol.* 13 (2026) 49, <http://dx.doi.org/10.1140/epjqt/s40507-026-00491-1>.
- [23] R. de Winter, B. van Stein, T. Bäck, SAMO-COBRA: a fast surrogate assisted constrained multi-objective optimization algorithm, in: *International Conference on Evolutionary Multi-Criterion Optimization*, 12654, Springer International Publishing, 2021, pp. 270–282, http://dx.doi.org/10.1007/978-3-030-72062-9_22.
- [24] L. Ruddigkeit, R. van Deursen, L.C. Blum, J.-L. Reymond, Enumeration of 166 billion organic small molecules in the chemical universe database GDB-17, *J. Chem. Inf. Model.* 52 (11) (2012) 2864–2875, <http://dx.doi.org/10.1021/ci300415d>.
- [25] A.S. Christensen, O.A. von Lilienfeld, Revised MD17 dataset (rMD17), 2020, <http://dx.doi.org/10.6084/m9.figshare.12672038.v3>, https://figshare.com/articles/dataset/Revised_MD17_dataset_rMD17/12672038. (Accessed April 2025).
- [26] A. Grisafi, D.M. Wilkins, M.J. Willatt, M. Ceriotti, Atomic-scale representation and statistical learning of tensorial properties, in: *Machine Learning in Chemistry: Data-Driven Algorithms, Learning Systems, and Predictions*, in: ACS Symposium Series, 1326, ACS Publications, 2019, pp. 1–21, <http://dx.doi.org/10.1021/bk-2019-1326.ch001>.
- [27] L. Zhang, J. Han, H. Wang, R. Car, W. E, Deep potential molecular dynamics: a scalable model with the accuracy of quantum mechanics, *Phys. Rev. Lett.* 120 (14) (2018) 143001, <http://dx.doi.org/10.1103/PhysRevLett.120.143001>.
- [28] J.S. Smith, O. Isayev, A.E. Roitberg, ANI-1: an extensible neural network potential with DFT accuracy at force field computational cost, *Chem. Sci.* 8 (4) (2017) 3192–3203, <http://dx.doi.org/10.1039/C6SC05720A>.
- [29] A. Musaelian, A. Johansson, S. Batzner, B. Kozinsky, Scaling the leading accuracy of deep equivariant models to biomolecular simulations of realistic size, in: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2023, pp. 1–12, <http://dx.doi.org/10.1145/3581784.3627041>.
- [30] X. Fu, A. Musaelian, A. Johansson, T. Jaakkola, B. Kozinsky, Learning interatomic potentials at multiple scales, 2023, arXiv preprint [arXiv:2310.13756](https://arxiv.org/abs/2310.13756).
- [31] L.-H. Lim, B.J. Nelson, What is... an equivariant neural network?, *Notices of the American Mathematical Society* 70 (4) (2023) 619–624, <http://dx.doi.org/10.1090/noti2666>.
- [32] T. Cohen, M. Welling, Group equivariant convolutional networks, in: *Proceedings of the 33rd International Conference on Machine Learning*, in: *Proceedings of Machine Learning Research*, 48, PMLR, 2016, pp. 2990–2999.
- [33] V. Garcia Satorras, E. Hoogeboom, M. Welling, E (n) equivariant graph neural networks, in: *Proceedings of the 38th International Conference on Machine Learning*, in: *Proceedings of Machine Learning Research*, 139, PMLR, 2021, pp. 9323–9332.
- [34] A. Navon, A. Shamsian, I. Achituve, E. Fetaya, G. Chechik, H. Maron, Equivariant architectures for learning in deep weight spaces, in: *Proceedings of the 40th International Conference on Machine Learning*, in: *Proceedings of Machine Learning Research*, 202, PMLR, 2023, pp. 25790–25816.
- [35] S.P. Ong, W.D. Richards, A. Jain, G. Hautier, M. Kocher, S. Cholia, D. Gunter, V.L. Chevrier, K.A. Persson, G. Ceder, Python materials genomics (pymatgen): A robust, open-source python library for materials analysis, *Comput. Mater. Sci.* 68 (2013) 314–319, <http://dx.doi.org/10.1016/j.commatsci.2012.10.028>.
- [36] A. Zur, T.C. McGill, Lattice match: An application to heteroepitaxy, *J. Appl. Phys.* 55 (2) (1984) 378–386, <http://dx.doi.org/10.1063/1.333084>.
- [37] W.G. Hoover, Canonical dynamics: Equilibrium phase-space distributions, *Phys. Rev. A* 31 (3) (1985) 1695–1697, <http://dx.doi.org/10.1103/PhysRevA.31.1695>.
- [38] S. Nosé, Constant Temperature Molecular Dynamics Methods, *Progr. Theoret. Phys. Suppl.* 103 (1991) 1–46, <http://dx.doi.org/10.1143/PTPS.103.1>.
- [39] S. Nosé, A unified formulation of the constant temperature molecular dynamics methods, *J. Chem. Phys.* 81 (1) (1984) 511–519, <http://dx.doi.org/10.1063/1.447334>.
- [40] G. Kresse, J. Furthmüller, Efficiency of ab-initio total energy calculations for metals and semiconductors using a plane-wave basis set, *Comput. Mater. Sci.* 6 (1) (1996) 15–50, [http://dx.doi.org/10.1016/0927-0256\(96\)00008-0](http://dx.doi.org/10.1016/0927-0256(96)00008-0).
- [41] G. Kresse, J. Furthmüller, Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set, *Phys. Rev. B* 54 (16) (1996) 11169–11186, <http://dx.doi.org/10.1103/PhysRevB.54.11169>.
- [42] G. Kresse, D. Joubert, From ultrasoft pseudopotentials to the projector augmented-wave method, *Phys. Rev. B* 59 (3) (1999) 1758–1775, <http://dx.doi.org/10.1103/PhysRevB.59.1758>.
- [43] J.P. Perdew, K. Burke, M. Ernzerhof, Generalized gradient approximation made simple, *Phys. Rev. Lett.* 77 (18) (1996) 3865–3868, <http://dx.doi.org/10.1103/PhysRevLett.77.3865>.
- [44] J.W. Furness, A.D. Kaplan, J. Ning, J.P. Perdew, J. Sun, Accurate and numerically efficient r2SCAN meta-generalized gradient approximation, *J. Phys. Chem. Lett.* 11 (19) (2020) 8208–8215, <http://dx.doi.org/10.1021/acs.jpcclett.0c02405>.
- [45] I. Batatia, D.P. Kovács, G.N.C. Simm, C. Ortner, G. Csányi, MACE: Higher order equivariant message passing neural networks for fast and accurate force fields, in: *Advances in Neural Information Processing Systems*, 35, Curran Associates, Inc., 2022, pp. 11423–11436.
- [46] I. Batatia, S. Batzner, D.P. Kovács, A. Musaelian, G.N.C. Simm, R. Drautz, C. Ortner, B. Kozinsky, G. Csányi, The design space of e(3)-equivariant atom-centred interatomic potentials, *Nat. Mach. Intell.* 7 (1) (2025) 56–67, <http://dx.doi.org/10.1038/s42256-024-00956-x>.
- [47] B. Coecke, R. Duncan, Interacting quantum observables, in: *Automata, Languages and Programming (ICALP 2008)*, in: *Lecture Notes in Computer Science*, 5126, Springer, 2008, pp. 298–310, http://dx.doi.org/10.1007/978-3-540-70583-3_25.
- [48] J. van de Wetering, ZX-calculus for the working quantum computer scientist, 2020, arXiv preprint [arXiv:2012.13966](https://arxiv.org/abs/2012.13966).
- [49] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, S. Woerner, The power of quantum neural networks, *Nat. Comput. Sci.* 1 (6) (2021) 403–409, <http://dx.doi.org/10.1038/s43588-021-00084-1>.
- [50] S.-i. Amari, Natural gradient works efficiently in learning, *Neural Comput.* 10 (2) (1998) 251–276, <http://dx.doi.org/10.1162/089976698300017746>.
- [51] J.R. McClean, S. Boixo, V.N. Smelyanskiy, R. Babbush, H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* 9 (1) (2018) 4812, <http://dx.doi.org/10.1038/s41467-018-07090-4>.
- [52] E. Luxenberg, S. Boyd, Exponentially weighted moving models, 2024, arXiv preprint [arXiv:2404.08136](https://arxiv.org/abs/2404.08136).
- [53] M. Larocca, N. Ju, D. García-Martín, P.J. Coles, M. Cerezo, Theory of over-parametrization in quantum neural networks, *Nat. Comput. Sci.* 3 (6) (2023) 542–551, <http://dx.doi.org/10.1038/s43588-023-00467-6>.
- [54] R. Ramakrishnan, P.O. Dral, M. Rupp, O.A. von Lilienfeld, Quantum chemistry structures and properties of 134 kilo molecules, *Sci. Data* 1 (2014) 140022, <http://dx.doi.org/10.1038/sdata.2014.22>.
- [55] Allegro+MLP: A modification of allegro to include MLP components, 2025, GitHub repository. <https://github.com/glatq/allegro>.
- [56] R. de Winter, SAMO-COBRA: A fast surrogate-assisted constrained multi-objective optimization algorithm, 2021, GitHub repository. <https://github.com/RoydeZomer/SAMO-COBRA>.