

LEARNING LONG-RANGE REPRESENTATIONS WITH EQUIVARIANT MESSAGES

Anonymous authors

Paper under double-blind review

ABSTRACT

Machine learning interatomic potentials trained on first-principles reference data are quickly becoming indispensable for computational physics, biology, and chemistry. Equivariant message-passing neural networks, including transformers, achieve state-of-the-art accuracy but rely on cutoff-based graphs, limiting their ability to capture long-range effects such as electrostatics or dispersion, as well as electron delocalization. While long-range correction schemes based on inverse power laws of interatomic distances have been proposed, they are unable to communicate higher-order geometric information and are thus limited in applicability and transferability. To address this shortcoming, we propose the use of equivariant, rather than scalar, charges for long-range interactions, and design a graph neural network architecture, LOREM, around this long-range message passing mechanism. We consider several datasets specifically designed to highlight non-local physical effects, and compare short-range message passing with different receptive fields to invariant and equivariant long-range message passing. Even though most approaches work for careful dataset-specific choices of their hyperparameters, LOREM works consistently without adjustments, with excellent benchmark performance.

1 INTRODUCTION

Machine learning interatomic potentials (MLIPs) for atomistic simulations are trained on quantum-mechanical simulations to predict energies and forces of new atomic structures. Most MLIPs assume locality: The energy of each atom depends only on neighbors within a cutoff radius. This leads to linear scaling with respect to the number of atoms but fails for systems with long-range interactions, such as electrostatics, dispersion, or electron delocalization. (Behler & Parrinello; Ko et al.; Tkatchenko & Scheffler; Grisafi & Ceriotti; Unke et al., b; Huguenin-Dumittan et al.) Several approaches aim to overcome this limitation. Message-passing graph neural networks (MPNNs, Gilmer et al.) overcome locality by iteratively exchanging information between neighboring atoms, but are still constrained by the number of message-passing steps, graph connectivity, and reduced information flow with increasing number of iterations (Cai & Wang; Alon & Yahav; Nigam et al.).

An alternative approach is to use physics-inspired corrections to the total energy, written as an inverse power law of interatomic distances, $1/r^p$, with $p = 1$ for charge–charge, $p = 3$ for dipole–dipole, and $p = 6$ for dispersion. Some models predict partial atomic charges, either using explicit charge labels and equilibration schemes, or learning them implicitly from energies and forces to include electrostatic terms (Unke et al., a; Ko et al.; Fedik et al.; Pellegrini et al.; Maruf et al.; Cheng, b). Physical long-range interactions can also serve as building block: For instance, Kosmala et al. propose Ewald message passing, i.e., the use of electrostatic interactions for message passing, and Grisafi & Ceriotti propose the long-distance equivariant (LODE) framework that uses inverse power law interactions to compute equivariant features.

In this work, we aim to combine the strengths of physical $1/r^p$ interactions, computational efficiency and physically meaningful asymptotic behavior, with the ability of equivariant message passing to communicate higher-order geometric information. Extending the idea of Ewald message passing, we treat charges as equivariant objects and use inverse power law potentials as a mechanism for long-range equivariant communication. Based on this mechanism, we design LOREM, an MLIP architecture that combines short- and long-range message passing. LOREM offers consistently high

accuracy across a series of long-range benchmark tasks, outperforming other short- and long-range message passing models.

Our contributions are:

- We introduce an equivariant, global, long-range message passing mechanism that leverages existing methods from computational physics for $O(N \log N)$ scaling in periodic systems,
- We design a novel MLIP architecture, LOREM, around this mechanism,
- We conduct a series of experiments to probe the limits of equivariant short-range message passing to model long-range physics.

2 BACKGROUND

Machine learning interatomic potentials Under the Born & Oppenheimer approximation, which decouples the nuclear and the electronic degrees of freedom, the atoms in a molecule or material move on a potential energy surface (PES) $E = E(\{(\mathbf{r}_i, Z_i) : i = 1 \dots N\})$ where $\mathbf{r}_i$ and Z_i are positions and atomic numbers for the N atoms; in a periodic system, i.e., crystals or liquids, the arrangement of N atoms is repeated periodically in space, described by a cell matrix $\mathbf{C}$. The forces, which drive the dynamics of the atoms, are defined as the derivatives $\mathbf{F}_i = -\nabla_{\mathbf{r}_i} E$. Traditionally, this PES has been approximated by physics-inspired analytical expressions, called force fields, that are parametrized manually or through global optimization. In the past decades, in tandem with the increasing availability of large datasets of quantum mechanical reference data, MLIPs have emerged as a less computationally efficient, but more accurate, data-driven alternative.

Atomistic graph neural networks Most MLIPs can be seen as graph neural networks (GNNs, Battaglia et al.) acting on a description of an arrangement of atoms in space as a geometric graph $\mathcal{G} = (\mathcal{E}, \mathcal{V})$ with edges $\mathcal{E}$ corresponding to interatomic relative-position vectors $\mathbf{r}_{ij}$ and nodes (or vertices) $\mathcal{V}$ corresponding to atoms. Edges are drawn between nodes that lie within a cutoff radius r_c of each other. $\mathbf{r}_{ij}$ are constructed to respect periodic boundary conditions. By restricting the range of interactions to neighbors on this graph, linear scaling with the number of atoms N (at constant density) can be achieved; efficient scaling with system size is required to make MLIPs practical for large-scale simulations. MLIPs typically predict E as a sum over atomic contributions, $E = \sum_{i=1}^N E_i$; the E_i are usually predicted based on the node features at each i .

Invariance and equivariance The potential energy E is invariant under translations, rotations, and permutations, i.e., relabeling, of atomic positions. MLIPs must respect these symmetries, at least approximately. Translation invariance is respected by construction in atomistic GNNs. Permutation invariance is typically ensured through commutative aggregation functions, for instance sums. Finally, rotation invariance can either be learned through data augmentation (Pozdnyakov & Ceriotti; Langer et al.), or ensured by requiring that internal features remain aware of their geometric meaning, i.e., that they are *equivariant* to rotations (Smidt). While in principle, MLIPs could be constructed from invariant features only, equivariant internal features have been found to improve accuracy and data efficiency by allowing the model to access orientation information (Thomas et al.; Batzner et al.; Batatia et al. (b)). One way to construct equivariant models is to associate features with an irreducible representation of $\text{SO}(3)$ or $\text{O}(3)$. We denote such objects of order l as tensors $\mathbf{X}^l$.

Long-range interactions The potential energy E arises from the many-body Schrödinger equation, which involves only Coulomb interactions between electrons and nuclei without any range separation. Efficient approximations, such as force fields or MLIPs, typically restrict interactions to local environments, motivated by the nearsightedness principle of electronic matter (Prodan & Kohn), which states that electronic properties are insensitive to distant perturbations. In the long-range regime, interactions reduce to inverse power laws $1/r^p$ of the interatomic distance. Since in nature no fixed nearsightedness length scale exists, MLIPs must capture both local many-body quantum effects, possibly extending beyond the model’s cutoff r_c , and formally infinite-range electrostatic interactions. Additional complexity arises from charge distributions that depend on distant atoms and from electron wavefunctions that may delocalize over large distances; see Section 6.

3 RELATED WORK

Equivariant message passing Bond-order potentials first introduced the idea of repeatedly updating atomic environments to extend interactions beyond r_c (Tersoff; Brenner), a principle now central to modern MLIPs. In MPNNs (Gilmer et al.), atoms exchange messages over M steps, so features and energies depend on neighbors within $M \cdot r_c$. Early models used invariant updates (Schütt et al., a; Xie & Grossman), later extended to equivariant ones (Gasteiger et al.; Schütt et al., b; Batatia et al., b; Batzner et al.; Frank et al., b). Recent universal MLIPs trained on big datasets (Batatia et al., a; Wood et al.; Mazitov et al.; Rhodes et al.) sometimes replace message passing with local self-attention. While effective at capturing semi-local interactions, MPNNs cannot model true long-range effects, since distant atoms without intermediates never interact, and many steps reduce expressivity (Cai & Wang; Alon & Yahav; Nigam et al.).

Physics-based long-range models Many long-range models explicitly add physics-inspired terms to E that capture interactions decaying more slowly with distance. A common example of such a form is pairwise additive: $E^{\text{long}} = \sum_{i < j} \phi(r_{ij}; q_i, q_j)$, where q_i are latent per atom descriptors (e.g., partial charges or polarizabilities) predicted by the model, and ϕ is a physically inspired, distance-based kernel. Usually, these q_i are predicted directly as scalar functions of local atomic features (Unke et al., a; Staacke et al.; Loche et al.; Kim et al.; Ji et al.; Kabylda et al.; Cheng, b). A modification of this approach is to include the q_i in an equilibration scheme, thus allowing these descriptors to capture information otherwise missed due to their initial dependence on local environments (Ko et al.; Pellegrini et al.; Maruf et al.). Grisafi & Ceriotti; Huguenin-Dumittan et al. propose to use physics-inspired kernels to compute long-range features, mathematically equivalent to a multipole expansion, but find that higher-order features contain little additional information.

Other long-range models Some models avoid handcrafted corrections and instead learn long-range interactions directly. Ewald message passing augments GNNs with Fourier-space invariants (Kosmala et al.), while fully connected approaches use all-to-all distances (Chmiela et al.) or global attention (Unke et al., a), though these lose efficiency or distance information. Linear-scaling attention with geometric embeddings (Frank et al., a) enables global orientation exchange but needs symmetrization. Alternatives include virtual nodes for global aggregation (Caruso et al.) or message passing in spherical harmonics space (Frank et al., c). Despite approximations, most methods still struggle to bridge periodic and non-periodic systems. Some methods outside atomistic modeling also aim to capture long-range effects (Dwivedi et al.; Bamberger et al.; Moskalov et al.; Zhdanov et al.).

4 EQUIVARIANT LONG-RANGE MESSAGE PASSING

Ewald summation, as introduced in Appendix C, allows the efficient, and convergent, evaluation of a potential at each atomic position, $V_i = V(\mathbf{r}_i) = \sum_{j=1}^N \sum_{\mathbf{n} \in \mathbb{Z}^3} q_j / |\mathbf{r}_i - (\mathbf{r}_j + \mathbf{C}\mathbf{n})|^p$. From the perspective of machine learning, this computation can be seen as a basic form of message passing, with radial filter $1/r_{ij}^p$, neighbor feature q_i , and resulting message V_i ; this correspondence was pointed out by Kosmala et al. and used by Grisafi & Ceriotti to compute general long-range features.

To bring it more in line with standard message passing, the operation can be carried out across an extra channel dimension, promoting a single feature q_i to a vector $\mathbf{q}_i$. This allows to communicate more information, but in this form, this information is restricted to geometric invariants. We propose to promote q_i to an *equivariant* tensor instead: $\mathbf{Q}_i^l$. This results in equivariant messages, or potentials:

$$\mathbf{v}_i^l = \sum_{j=1}^N \sum_{\mathbf{n} \in \mathbb{Z}^3} \frac{\mathbf{Q}_j^l}{|\mathbf{r}_i - (\mathbf{r}_j + \mathbf{C}\mathbf{n})|^p}. \quad (1)$$

To see that the result is equivariant, recall that adding two equivariant objects yields another equivariant object, and that multiplication with a prefactor, provided that the factor is shared across all entries for a given l , also retains equivariance. It is then easy to see that multiplying $\mathbf{Q}_i^l$ with a factor $1/r_{ij}^p$ yields an equivariant, and that the sum in Eq. (1) also yields an equivariant, since all summands are equivariant objects. We argue and numerically confirm in Appendix A that a compensating background charge correction preserves equivariance.

Overall, this approach allows the global aggregation of equivariant messages in a way that is amenable to efficient implementation, scaling $O(N \log N)$ in periodic systems and – at least in principle, see Appendix C – $O(N)$ in finite systems. The method also naturally accommodates the asymptotic power-law decay intrinsic to most long-range effects in physical systems.

5 LOREM




Figure 1: Sketch of the LOREM architecture.

LOREM follows the blueprint of equivariant MPNNs, making a few modifications: Scalar and spherical features are handled separately, following Frank et al. (c;b), and mixed using a variant of the power spectrum (Bartók et al.) and a residual update block (He et al.).

Overview LOREM, illustrated in Fig. 1 maps a point cloud of N atomic positions $\{r_i\}$, potentially contained in a periodic unit cell C , and labeled with chemical species $\{Z_i\}$ to atomic energies $\{E_i\}$. Internally, the point cloud is processed as a graph, with connectivity between nodes, the atoms, defined by a cutoff radius r_c . Accordingly, edges correspond to vectors $r_{ij} = r_j - r_i$. Node features are updated through either **short-range message passing** or **long-range message passing**. Spherical information is used to update scalar features by computing their **spherical norm** and passing it through an **update** block. Updates of scalar features are followed by a residual prediction of atomic energy contributions. Selected blocks are discussed below; the remaining in Appendix E.

Short-range message passing Distances r_{ij} are expanded in Bernstein polynomials multiplied with f_{cut} , the cosine cutoff function, yielding ρ_{ij} . Orientations r_{ij} are expanded in spherical harmonics Y_{ij}^l . Initial scalar features are a learned embedding of chemical species. At each short-range message

passing step,¹ edge features $\mathbf{K}_{ij}$ are obtained by concatenating ${}^k\mathbf{P}_i$ and ${}^k\mathbf{P}_j$. They are the input for a learned linear transformation of the radial basis; the result is summed to update node features. Linearly transformed edge features are also used as prefactors for $\mathbf{Y}_{ij}^l$, which participate in a tensor product with ${}^k\mathbf{S}_j^l$; the results are aggregated into a spherical message, which updates the spherical features via a tensor product, resulting in ${}^{k+1}\mathbf{S}_i^l$. The spherical norm of the updated spherical node features is then used to further update the scalar features, finally yielding ${}^{k+1}\mathbf{P}_i$. This process is repeated M times (the total number of short-range message passing steps).

Long-range message passing To minimize computational cost, which is proportional to the number of channels over which Ewald summation is carried out, node features are transformed into low-dimensional charges: The scalar features ${}^M\mathbf{P}_i$ are transformed into a single charge q_i , and spherical features ${}^M\mathbf{S}_i^l$ are likewise, using a linear transformation followed by a self tensor product, transformed into $\mathbf{Q}_i^l$, typically with a reduced maximum spherical harmonics degree and singular feature dimension. Unless otherwise noted, we use $l_{\max} = 2$, which yields a total of 10 charge channels.² After Ewald summation, which is carried out in parallel across charge channels, the potentials are split into the purely scalar V_i and the spherical $\mathbf{V}_i^l$. The spherical potentials are combined with spherical node features through a tensor product; the spherical norm of the result is concatenated with the scalar potential to update scalar representations. We find that using only $p=1$, i.e., Coulomb interaction, rather than a set of different exponents, is sufficient in practice.

6 EXPERIMENTS

We perform experiments on a number of existing benchmark tasks designed to probe the ability of MLIPs to model long-range interactions, comparing LOREM to short and long-range MLIPs. The experiments are divided into two parts: In Section 6.1, we compare LOREM with other models using standard settings. We find that LOREM performs well, but also observe that most tasks can be solved by models that do not consider long-range interactions at all – the effective interaction range of typical message passing models is sufficient. However, predictions break down beyond this interaction range. In Section 6.2, we probe this limit of short-range message passing. We find that while careful consideration of hyperparameters is required to resolve long-range interactions with short-range models, LOREM can solve these benchmarks without adaptation.

Models We compare LOREM with a number of purely short-ranged MLIPs: MACE and PET, as well as the recently introduced CACE-LES model that combines short-range message passing with a scalar long-range part, and 4G-NN, which includes a physics-based long-range energy contribution and charge equilibration. Full details on model descriptions and training can be found in Appendix F.

The datasets and associated benchmark tasks used in our experiments are briefly introduced below; additional details are given in Appendix B. An overview of validation or test set (where available) metrics for all datasets is given in Table 1. LOREM performs very well across datasets, and is competitive with, or more accurate than, other models.

MgO surface This benchmark task is the first in a series designed by Ko et al. to highlight the need for long-range information in MLIPs. Illustrated in Fig. 2A, it consists of a magnesium oxide (MgO) surface on which a gold (Au_2) dimer is placed. Depending on the presence of an aluminum (Al) dopant deep inside the surface, the lowest-energy position of the gold dimer is either in the ‘wetting’ (flat), or ‘non-wetting’ (standing up), position. The benchmark consists of two parts: Correctly identifying the ordering between wetting and non-wetting geometries in the doped and undoped case, and reproducing the energy-distance curve for the non-wetting geometry, in particular the local minimum corresponding to the equilibrium distance.

¹In the initial message passing step, no spherical node features are available, and therefore the tensor product with $\mathbf{Y}_{ij}^l$ is omitted and initial spherical node features are obtained via a self tensor product (dotted lines).

²Two with $l=0$: one from the scalar features and one from the scalar component of the spherical features.

³4G-NN training requires DFT-computed charges, which are not available for the biodimers and cumulene.

Table 1: Root mean squared errors for energy E and forces F for datasets and models used in Section 6.1. Where available, a held-out test set was used; for small datasets, the validation set was used instead and indicated in the table. The lowest error is indicated in bold. The second line after each model name indicates the number of short-range (SR) message-passing steps and whether some form of long-range (LR) interactions are included in the model.³

Dataset		LOREM 1×SR + LR	CACE-LES 1×SR + LR	MACE 2×SR	PET 2×SR	4G-NN 1×SR + LR
MgO surface (Validation set)	E (meV/at)	0.064	0.071	0.376	0.210	0.219
	F (meV/Å)	4.076	7.913	5.971	6.261	66.000
NaCl cluster (Validation set)	E (meV/at)	0.112	0.210	1.681	1.517	0.481
	F (meV/Å)	1.155	9.784	40.219	42.438	32.780
Biodimers	E (meV/at)	0.222	2.259	7.793	6.758	–
	F (meV/Å)	1.646	3.163	16.150	16.470	–
Cumulene	E (meV/at)	3.309	17.803	12.592	3.205	–
	F (meV/Å)	50.084	147.616	104.318	46.905	–

NaCl cluster Similar to the presence of a dopant modifying the potential-energy surface for the gold dimer in the MgO surface task, this benchmark by Ko et al. relies on the presence or absence of a sodium atom at one end of a charged sodium chloride (NaCl) cluster changing the behavior of a sodium atom at the opposite end: Since one atom is removed while the charge remains constant, the charge must redistribute over the remaining atoms. The benchmark task consists of reproducing the location of the local minimum and the energy profile when moving the sodium atom farthest from the removed one, indicated in Fig. 2C.

Cumulene The cumulene benchmark task was proposed by Unke et al. (b) as an example of a long-range problem that is due to a non-local effect of the electronic structure of a molecule. This molecule consists of a chain of nine carbon atoms, with a pair of hydrogen atoms, the rotors, at the opposite ends. The orientation of the rotors determines the shape of the atomic orbitals for each carbon, which propagates along the chain to the other end; in the absence of bending or stretching, the energy is therefore fully determined by the relative orientation of the rotors. The task is to recover the energy profile of this rotation for an idealized fully extended geometry, illustrated in Fig. 3.

Biodimers This benchmark dataset by Huguenin-Dumittan et al. consists of pairs of relaxed organic molecules placed at distances of 4 Å to 15 Å from each other, illustrated in Fig. 4A. Depending on the chemical nature of the molecules, interactions between the molecules take different asymptotic power law forms, ranging from charge-charge $1/r$ interactions to apolar-apolar $1/r^6$ interactions. Here, we simply compare energy and force prediction errors on test sets stratified by the dominating power-law interaction.

S_N2 reactions Finally, the S_N2 reactions benchmark was introduced by Frank et al. (a) to probe the ability of MLIPs to model the long-range interactions required to mediate gas-phase chemical reactions. It consists of the nucleophilic substitution of methyl halides by another halide ion: $X^- + H_3C-Y \rightarrow X-CH_3 + Y^-$ where X,Y = F, Cl, Br, or I. The benchmark task in this case is predicting the energy along the reaction coordinate, i.e., correctly modeling the potential energy profile as the two reactants approach one another, react, and separate again.

6.1 STANDARD SETTINGS

We compare LOREM to other models using their standard settings (see Appendix F).

MgO surface The results for this experiment can be seen in Fig. 2B: Despite being designed to require long-range interactions, this benchmark task can be solved by both purely short-range and long-range message-passing models. In this case, the success of short-range message passing is due to the small size of this benchmark system: With an effective cutoff radius above 6 Å, centrally




Figure 2: (A) Au_2 dimer on MgO surface, showing both wetting and non-wetting geometries, as well as the Al dopant. (B) Energy over distance d for the non-wetting geometry for the doped and undoped surface. The minima are indicated with a diamond symbol; the reference energy curve is drawn in grey. Offsets are added to distinguish the curves and the value at the minimum is subtracted. (C) Na_9Cl_8^+ (top) and Na_8Cl_8^+ (bottom) cluster, the moving atom is marked with transparent copies of itself, and the distance of interest is labeled with d . (D) Energy over distance for both clusters.

located atoms can ‘see’ the full system and hence, MACE and PET can solve this task. All models also resolve the orientation preference for the Au_2 dimer between the doped and undoped surfaces.

NaCl cluster The results of this experiment, seen in Fig. 2D, are drastically different from the previous one: Here, only models with a long-range component are able to resolve the difference between Na_9Cl_8^+ and Na_8Cl_8^+ . All such models show excellent agreement with reference values. The failure of short-range message-passing is due to the larger system size compared to the MgO surface: Here, effective cutoff radii exceeding $10.5 \text{ \AA}$ are required to solve the benchmark task.

Cumulene Resolving the cumulene energy profile, seen in Fig. 3, requires simultaneous knowledge of the orientation of both rotors at opposite ends of the molecule. This can be achieved in two ways: Through equivariant short-range message passing (MACE, PET), provided that the chain is not too long (see Section 6.2), or through *equivariant* long-range message passing (LOREM). For this reason, CACE-LES cannot solve this benchmark: Scalar charges are not sufficiently expressive to communicate relative orientation. All equivariant models that are able to access this information can solve this benchmark, achieving good agreement with the reference data. PET resolves the dihedral angle as well, but requires long training and an adaptation in the number of transformer layers to succeed at this benchmark (see Appendix F). We note that the sharp cusp at 180° is an artifact of the underlying reference method; it is a desirable behavior of MLIPs to smoothen it out.




Figure 3: (A) Illustration of cumulene laid flat, indicating relevant distances between atoms. (B) Energy profile over a 90° rotation of one rotor. The minimum value of each curve is subtracted before plotting. The inset shows a 3D representation of cumulene, defining the dihedral angle θ .




Figure 4: Mean absolute error on forces for different models on the different dimer classes: Apolar-apolar (AA), charge-apolar (CA), charge-charge (CC), charge-polar (CP), polar-apolar (PA), polar-polar (PP). A charge-charge pair is shown as inset.

Biodimers Since the pairs of molecules in the biodimers benchmark are placed at separations up to $15 \text{ \AA}$, much beyond the cutoff used for graph construction, high accuracy requires a long-range component. This is confirmed by the results of this experiment, seen in Fig. 4: The long range LOREM and CACE-LES models yield lower error than MACE and PET. Predictive error varies between dimer classes, i.e., the expected type of inverse power law interaction, for all models, with the exception for LOREM, which yields consistent, and in all but one class the highest, accuracy. This may be due to the addition of a nonlinearity after long-range message passing, which allows the model to correct the mismatch between the power law used for message passing, $1/r$, and the different exponents in the dimer classes.

6.2 LIMITS OF SHORT-RANGE MESSAGE PASSING

In the previous experiments, we observed that the performance of models with short-range message passing strongly depends on the match between the effective interaction cutoff and the problem to be solved. The case of biodimers, demonstrates clearly that message passing cannot resolve interactions where no intermediate atoms are present. To study these cases, we perform additional experiments using LOREM, with and without long-range message passing, and with different cutoffs.

Cumulene with different cutoffs We probe the dependence of the performance of message-passing models on their hyperparameters by training a set of LOREM models with and without long-range message passing at different cutoffs ($2.5 \text{ \AA}$, $3 \text{ \AA}$ and $3.5 \text{ \AA}$) and numbers of short-range message passing steps (1, 2, 3 and 4). Small cutoffs are chosen to simulate the longer chain lengths of real biomolecular systems. The results are shown in Table 2 of Appendix D: In all cases, models that include long-range message passing are able to resolve the angle. Short-range models, on the other hand, can only resolve the angle in certain combinations of hyperparameters: One message passing step is never sufficient. For two, a minimum cutoff of $3.5 \text{ \AA}$ is required. For three message passing steps, $3.0 \text{ \AA}$ is required. Therefore, an effective cutoff substantially larger than the $5.8 \text{ \AA}$ distance

between the central carbon atom in the chain and the hydrogen rotors is needed. This is because MPNNs rely on the input graph structure (see Fig. 3A) for information flow: Below $r_c = 2.6$ Å, connections in the graph only extend to nearest neighbors, with the exception of the rotors and the second-to-last carbon atoms. Consequently, at least four message passing steps are required. This example illustrates that short-range message passing is difficult to apply to this class of problems: Parameters have to be adapted to the chain length, with careful consideration of graph structure. In contrast, long-range message passing is robust, requiring no change in parameters to solve this task.




Figure 5: Energy over the reaction coordinate for the nucleophilic substitution reaction $\text{Cl}^- + \text{H}_3\text{C}-\text{Br} \rightarrow \text{Cl}-\text{CH}_3 + \text{Br}^-$; snapshots are shown as insets.

$\text{S}_{\text{N}}2$ reactions Solving the benchmark task for the $\text{S}_{\text{N}}2$ reactions dataset requires long-range interactions, since it involves modeling intra-molecular interactions over distances exceeding the typical cutoffs used for MLIPs. This is confirmed by Fig. 5, which compares the performance of LOREM with and without long-range message passing. The former can reproduce the energy over the course of the reaction with excellent accuracy, including the tails as the reactants approach and separate. The latter, which does not include long-range interactions, cannot account for the tails and consequently predicts a constant once the molecules separate more than 5 Å.

7 DISCUSSION

We introduced a simple yet effective equivariant long-range message passing scheme: Latent equivariant charges are predicted from local node features, and well-established techniques for evaluating inverse power-law potentials are used to efficiently compute long-range messages. Building on this, we developed LOREM, a MLIP architecture that achieves consistently strong performance across benchmarks that require accurate modeling of long-range interactions.

By construction, our model assumes that interactions decay asymptotically with distance. While it is therefore not suited for learning truly global representations that have no notion of locality, this limitation is largely theoretical for physical systems: Electrostatics dominates most long-range behavior in realistic systems, and even other effects typically do not extend over arbitrary distances. A more practical limitation arises in non-periodic systems, where the cost of long-range message evaluation scales as $O(N^2)$. Although this is acceptable for small molecules, this may become a bottleneck for larger systems. Linear scaling can be achieved using fast multipole methods (Greengard & Rokhlin; Andy L Jones) or alternative methods like Euclidean Fast Attention (Frank et al., a).

We also investigated the capabilities of purely short-range message passing and found that, in many cases, it performs very well—even on datasets explicitly designed to require long-range interactions or charge equilibration. However, its success depends on matching the cutoff radius and number of message passing steps to the specific task, a process that can be both tedious and error-prone. More fundamentally, short-range methods cannot resolve interactions between distant atoms without intermediaries. Our augmentation with long-range message passing overcomes these limitations.

Finally, our results underscore a broader issue: the lack of challenging long-range benchmarks. Many current datasets are based on simplified, small-scale systems and fail to capture the complexity of real-world applications where long-range interactions are essential. Addressing this gap is a key direction for future work.

REPRODUCIBILITY STATEMENT

Code, data, configuration files, and trained models are available at doi:10.5281/zenodo.17194573. Scripts include data pre-processing, model training, model evaluation, and the creation of figures and tables in this work. Hyperparameters are also described in Appendix F.

REFERENCES

- Joseph W. Abbott, Carlos Mera Acosta, Alaa Akkoush, Alberto Ambrosetti, Viktor Atalla, Alexej Bagrets, Jörg Behler, Daniel Berger, Björn Bieniek, Jonas Björk, Volker Blum, Saeed Bohloul, Connor L. Box, Nicholas Boyer, Danilo Simoes Brambila, Gabriel A. Bramley, Kyle R. Bryenton, María Camarasa-Gómez, Christian Carbogno, Fabio Caruso, Sucismita Chutia, Michele Ceriotti, Gábor Csányi, William Dawson, Francisco A. Delesma, Fabio Della Sala, Bernard Delley, Robert A. DiStasio Jr, Maria Dragoumi, Sander Driessen, Marc Dvorak, Simon Erker, Ferdinand Evers, Eduardo Fabiano, Matthew R. Farrow, Florian Fiebig, Jakob Filser, Lucas Foppa, Lukas Gallandi, Alberto Garcia, Ralf Gehrke, Simiam Ghan, Luca M. Ghiringhelli, Mark Glass, Stefan Goedecker, Dorothea Golze, Matthias Gramzow, James A. Green, Andrea Grisafi, Andreas Grüneis, Jan Günzl, Stefan Gutzeit, Samuel J. Hall, Felix Hanke, Ville Havu, Xingtao He, Joscha Hekele, Olle Hellman, Uthpala Herath, Jan Hermann, Daniel Hernangómez-Pérez, Oliver T. Hofmann, Johannes Hoja, Simon Hollweger, Lukas Hörmann, Ben Hourahine, Wei Bin How, William P. Huhn, Marcel Hülsberg, Timo Jacob, Sara Panahian Jand, Hong Jiang, Erin R. Johnson, Werner Jürgens, J. Matthias Kahk, Yosuke Kanai, Kisung Kang, Petr Karpov, Elisabeth Keller, Roman Kempf, Danish Khan, Matthias Kick, Benedikt P. Klein, Jan Kloppenburg, Alexander Knoll, Florian Knoop, Franz Knuth, Simone S. Köcher, Jannis Kockläuner, Sebastian Kokott, Thomas Körzdörfer, Hagen-Henrik Kowalski, Peter Kratzer, Pavel K° us, Raul Laasner, Bruno Lang, Björn Lange, Marcel F. Langer, Ask Hjorth Larsen, Hermann Lederer, Susi Lehtola, Maja-Olivia Lenz-Himmer, Moritz Leucke, Sergey Levchenko, Alan Lewis, O. Anatole von Lilienfeld, Konstantin Lion, Werner Lipsunen, Johannes Lischner, Yair Litman, Chi Liu, Qing-Long Liu, Andrew J. Logsdail, Michael Lorke, Zekun Lou, Iuliia Mandzhieva, Andreas Marek, Johannes T. Margraf, Reinhard J. Maurer, Tobias Melson, Florian Merz, Jörg Meyer, Georg S. Michelitsch, Teruyasu Mizoguchi, Evgeny Moerman, Dylan Morgan, Jack Morgenstein, Jonathan Moussa, Akhil S. Nair, Lydia Nemec, Harald Oberhofer, Alberto Otero-de-la Roza, Ramón L. Panadés-Barrueta, Thanush Patlolla, Mariia Pogodaeva, Alexander Pöpl, Alastair J. A. Price, Thomas A. R. Purcell, Jingkai Quan, Nathaniel Raimbault, Markus Rampp, Karsten Rasim, Ronald Redmer, Xinguo Ren, Karsten Reuter, Norina A. Richter, Stefan Ringe, Patrick Rinke, Simon P. Rittmeyer, Herzain I. Rivera-Arrieta, Matti Ropo, Mariana Rossi, Victor Ruiz, Nikita Rybin, Andrea Sanfilippo, Matthias Scheffler, Christoph Scheurer, Christoph Schober, Franziska Schubert, Tonghao Shen, Christopher Shepard, Honghui Shang, Kiyoo Shibata, Andrei Sobolev, Ruyi Song, Aloysius Soon, Daniel T. Speckhard, Pavel V. Stishenko, Muhammad Tahir, Izumi Takahara, Jun Tang, Zechen Tang, Thomas Theiss, Franziska Theiss, Alexandre Tkatchenko, Milica Todorović, George Trenins, Oliver T. Unke, Álvaro Vázquez-Mayagoitia, Oscar van Vuren, Daniel Waldschmidt, Han Wang, Yanyong Wang, Jürgen Wieferink, Jan Wilhelm, Scott Woodley, Jianhang Xu, Yong Xu, Yi Yao, Yingyu Yao, Mina Yoon, Victor Wen-zhe Yu, Zhenkun Yuan, Marios Zacharias, Igor Ying Zhang, Min-Ye Zhang, Wentao Zhang, Rundong Zhao, Shuo Zhao, Ruiyi Zhou, Yuanyuan Zhou, and Tong Zhu. Roadmap on Advancements of the FHI-aims Software Package.
- Uri Alon and Eran Yahav. On the Bottleneck of Graph Neural Networks and its Practical Implications. URL <https://openreview.net/forum?id=i80OPhOCVH2>.
- Andy L Jones. pybbfmm. URL <https://www.github.com/andyljones/pybbfmm>.
- Jacob Bamberger, Benjamin Gutteridge, Scott le Roux, Michael M. Bronstein, and Xiaowen Dong. On Measuring Long-Range Interactions in Graph Neural Networks. URL [https://openreview.net/forum?id=2fBcAOi8l0&referrer=%5Bthe%20profile%20of%20Jacob%20Bamberger%5D\(%2Fprofile%3Fid%3D~Jacob_Bamberger1\)](https://openreview.net/forum?id=2fBcAOi8l0&referrer=%5Bthe%20profile%20of%20Jacob%20Bamberger%5D(%2Fprofile%3Fid%3D~Jacob_Bamberger1)).
- Albert P. Bartók, Risi Kondor, and Gábor Csányi. On representing chemical environments. 87(18): 184115. doi: 10.1103/PhysRevB.87.184115.

- Ilyes Batatia, Philipp Benner, Yuan Chiang, Alin M. Elena, Dávid P. Kovács, Janosh Riebesell, Xavier R. Advincula, Mark Asta, Matthew Avaylon, William J. Baldwin, Fabian Berger, Noam Bernstein, Arghya Bhowmik, Samuel M. Blau, Vlad C̃ arare, James P. Darby, Sandip De, Flaviano Della Pia, Volker L. Deringer, Rokas Elijõ sius, Zakariya El-Machachi, Fabio Falcioni, Edwin Fako, Andrea C. Ferrari, Annalena Genreith-Schriever, Janine George, Rhys E. A. Goodall, Clare P. Grey, Petr Grigorev, Shuang Han, Will Handley, Hendrik H. Heenen, Kersti Hermansson, Christian Holm, Jad Jaafar, Stephan Hofmann, Konstantin S. Jakob, Hyunwook Jung, Venkat Kapil, Aaron D. Kaplan, Nima Karimitari, James R. Kermode, Namu Kroupa, Jolla Kullgren, Matthew C. Kuner, Domantas Kuryla, Guoda Liepuoniute, Johannes T. Margraf, Ioan-Bogdan Magd̃ au, Angelos Michaelides, J. Harry Moore, Aakash A. Naik, Samuel P. Niblett, Sam Walton Norwood, Niamh O'Neill, Christoph Ortner, Kristin A. Persson, Karsten Reuter, Andrew S. Rosen, Lars L. Schaaf, Christoph Schran, Benjamin X. Shi, Eric Sivonxay, Tamás K. Stenczel, Viktor Svahn, Christopher Sutton, Thomas D. Swinburne, Jules Tilly, Cas van der Oord, Eszter Varga-Umbrich, Tejs Vegge, Martin Vondrák, Yangshuai Wang, William C. Witt, Fabian Zills, and Gábor Csányi. A foundation model for atomistic materials chemistry, a.
- Ilyes Batatia, David P. Kovacs, Gregor Simm, Christoph Ortner, and Gabor Csanyi. MACE: Higher Order Equivariant Message Passing Neural Networks for Fast and Accurate Force Fields. 35:11423–11436, b. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/4a36c3c51af11ed9f34615b81edb5bbc-Abstract-Conference.html.
- Peter W. Battaglia, Jessica B. Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, Caglar Gulcehre, Francis Song, Andrew Ballard, Justin Gilmer, George Dahl, Ashish Vaswani, Kelsey Allen, Charles Nash, Victoria Langston, Chris Dyer, Nicolas Heess, Daan Wierstra, Pushmeet Kohli, Matt Botvinick, Oriol Vinyals, Yujia Li, and Razvan Pascanu. Relational inductive biases, deep learning, and graph networks.
- Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P. Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E. Smidt, and Boris Kozinsky. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. 13(1). ISSN 2041-1723. doi: 10.1038/s41467-022-29939-5.
- Jörg Behler and Michele Parrinello. Generalized Neural-Network Representation of High-Dimensional Potential-Energy Surfaces. 98(14):146401. ISSN 0031-9007, 1079-7114. doi: 10.1103/physrevlett.98.146401.
- M. Born and R. Oppenheimer. Zur Quantentheorie der Molekeln. 389(20):457–484. ISSN 0003-3804, 1521-3889. doi: 10.1002/andp.19273892002.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Neca, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: Composable transformations of Python+NumPy programs. URL <http://github.com/google/jax>.
- Donald W. Brenner. Empirical potential for hydrocarbons for use in simulating the chemical vapor deposition of diamond films. 42(15):9458–9471. doi: 10.1103/PhysRevB.42.9458.
- Lori A Burns, John C Faver, Zheng Zheng, Michael S Marshall, Daniel GA Smith, Kenno Vanommeslaeghe, Alexander D MacKerell, Kenneth M Merz, and C David Sherrill. The biofragment database (bfdb): An open-data platform for computational chemistry analysis of noncovalent interactions. *The Journal of chemical physics*, 147(16), 2017.
- Chen Cai and Yusu Wang. A Note on Over-Smoothing for Graph Neural Networks.
- Alessandro Caruso, Jacopo Venturin, Lorenzo Giambagli, Edoardo Rolando, Frank Noé, and Cecilia Clementi. Extending the RANGE of Graph Neural Networks: Relaying Attention Nodes for Global Encoding.
- Bingqing Cheng. Cartesian atomic cluster expansion for machine learning interatomic potentials. 10(1):157, a. ISSN 2057-3960. doi: 10.1038/s41524-024-01332-4.

- Bingqing Cheng. Latent Ewald summation for machine learning of long-range interactions. 11(1):80, b. ISSN 2057-3960. doi: 10.1038/s41524-025-01577-7.
- Stefan Chmiela, Huziel E. Sauceda, Klaus-Robert Müller, and Alexandre Tkatchenko. Towards exact molecular dynamics simulations with machine-learned force fields. 9(1):3887. ISSN 2041-1723. doi: 10.1038/s41467-018-06169-2.
- Tom Darden, Darrin York, and Lee Pedersen. Particle mesh Ewald: An $N \log(N)$ method for Ewald sums in large systems. 98(12):10089–10092. ISSN 0021-9606. doi: 10.1063/1.464397.
- Vijay Prakash Dwivedi, Ladislav Rampásek, Mikhail Galkin, Ali Parviz, Guy Wolf, Anh Tuan Luu, and Dominique Beaini. Long Range Graph Benchmark. URL <https://openreview.net/forum?id=in7XC5RcjEn>.
- P. P. Ewald. Die Berechnung optischer und elektrostatischer Gitterpotentiale. 369(3):253–287. ISSN 1521-3889. doi: 10.1002/andp.19213690304.
- Nikita Fedik, Roman Zubatyuk, Maksim Kulichenko, Nicholas Lubbers, Justin S. Smith, Benjamin Nebgen, Richard Messerly, Ying Wai Li, Alexander I. Boldyrev, Kipton Barros, Olexandr Isayev, and Sergei Tretiak. Extending machine learning beyond interatomic potentials for predicting molecular properties. 6(9):653–672. ISSN 2397-3358. doi: 10.1038/s41570-022-00416-3.
- J. Thorben Frank, Stefan Chmiela, Klaus-Robert Müller, and Oliver T. Unke. Euclidean Fast Attention: Machine Learning Global Atomic Representations at Linear Cost, a.
- J. Thorben Frank, Oliver T. Unke, Klaus-Robert Müller, and Stefan Chmiela. A Euclidean transformer for fast and stable machine learned force fields. 15(1):6539, b. ISSN 2041-1723. doi: 10.1038/s41467-024-50620-6.
- Thorben Frank, Oliver Unke, and Klaus-Robert Müller. So3krates: Equivariant attention for interactions on arbitrary length-scales in molecular systems. 35:29400–29413, c. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/bcf4ca90a8d405201d29dd47d75ac896-Abstract-Conference.html.
- Johannes Gasteiger, Janek Groß, and Stephan Günnemann. Directional Message Passing for Molecular Graphs. URL <https://openreview.net/forum?id=BleWbxStPH>.
- Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural Message Passing for Quantum Chemistry. In *Proceedings of the 34th International Conference on Machine Learning*, pp. 1263–1272. PMLR. URL <https://proceedings.mlr.press/v70/gilmer17a.html>.
- L. Greengard and V Rokhlin. A fast algorithm for particle simulations. 73(2):325–348. ISSN 0021-9991. doi: 10.1016/0021-9991(87)90140-9.
- Andrea Grisafi and Michele Ceriotti. Incorporating long-range physics in atomic-scale machine learning. 151(20):204105. ISSN 0021-9606. doi: 10.1063/1.5128375.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition.
- R. W. Hockney and J. W. Eastwood. *Computer Simulation Using Particles*. CRC Press. ISBN 978-0-367-80693-4. doi: 10.1201/9780367806934.
- Kevin K. Huguenin-Dumittan, Philip Loche, Ni Haoran, and Michele Ceriotti. Physics-Inspired Equivariant Descriptors of Nonbonded Interactions. 14(43):9612–9618. doi: 10.1021/acs.jpcclett.3c02375.
- Yajie Ji, Jiuyang Liang, and Zhenli Xu. Machine-Learning Interatomic Potentials for Long-Range Systems.
- Adil Kabylda, J. Thorben Frank, Sergio Suarez Dou, Almaz Khabibrakhmanov, Leonardo Medrano Sandonas, Oliver T. Unke, Stefan Chmiela, Klaus-Robert Müller, and Alexandre Tkatchenko. Molecular Simulations with a Pretrained Neural Network and Universal Pairwise Force Fields.

- Dongjin Kim, Daniel S. King, Peichen Zhong, and Bingqing Cheng. Learning charges and long-range interactions from energies and forces.
- Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization.
- Tsz Wai Ko, Jonas A. Finkler, Stefan Goedecker, and Jörg Behler. A fourth-generation high-dimensional neural network potential with accurate electrostatics including non-local charge transfer. 12(1):398. ISSN 2041-1723. doi: 10.1038/s41467-020-20427-2.
- Arthur Kosmala, Johannes Gasteiger, Nicholas Gao, and Stephan Günnemann. Ewald-based Long-Range Message Passing for Molecular Graphs. In *Proceedings of the 40th International Conference on Machine Learning*, pp. 17544–17563. PMLR. URL <https://proceedings.mlr.press/v202/kosmala23a.html>.
- Marcel F Langer, Sergey N Pozdnyakov, and Michele Ceriotti. Probing the effects of broken symmetries in machine learning. 5(4):04LT01. ISSN 2632-2153. doi: 10.1088/2632-2153/ad86a0.
- Philip Loche, Kevin K. Huguenin-Dumittan, Melika Honarmand, Qianjun Xu, Egor Rumiantsev, Wei Bin How, Marcel F. Langer, and Michele Ceriotti. Fast and flexible long-range models for atomistic machine learning. 162(14):142501. ISSN 0021-9606. doi: 10.1063/5.0251713.
- Moin Uddin Maruf, Sungmin Kim, and Zeeshan Ahmad. Equivariant Machine Learning Interatomic Potentials with Global Charge Redistribution.
- Arslan Mazitov, Filippo Bigi, Matthias Kellner, Paolo Pegolo, Davide Tisi, Guillaume Fraux, Sergey Pozdnyakov, Philip Loche, and Michele Ceriotti. PET-MAD, a lightweight universal interatomic potential for advanced materials modeling.
- Artem Moskalev, Mangal Prakash, Junjie Xu, Tianyu Cui, Rui Liao, and Tommaso Mansi. Geometric Hyena Networks for Large-scale Equivariant Learning. URL [https://openreview.net/forum?id=jJRkkPr474&referrer=%5Bthe%20profile%20of%20Mangal%20Prakash%5D\(%2Fprofile%3Fid%3D~Mangal_Prakash1\)](https://openreview.net/forum?id=jJRkkPr474&referrer=%5Bthe%20profile%20of%20Mangal%20Prakash%5D(%2Fprofile%3Fid%3D~Mangal_Prakash1)).
- Jigyasa Nigam, Sergey Pozdnyakov, Guillaume Fraux, and Michele Ceriotti. Unified theory of atom-centered representations and message-passing machine-learning schemes. 156(20):204115. ISSN 0021-9606. doi: 10.1063/5.0087042.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc. URL https://papers.nips.cc/paper_files/paper/2019/hash/bdbca288fee7f92f2bfa9f7012727740-Abstract.html.
- Franco Pellegrini, Ruggero Lot, Yusuf Shaidu, and Emine Küçükbenli. PANNA 2.0: Efficient neural network interatomic potentials and new architectures. 159(8):084117. ISSN 0021-9606. doi: 10.1063/5.0158075.
- John P. Perdew, Kieron Burke, and Matthias Ernzerhof. Generalized Gradient Approximation Made Simple. 77(18):3865–3868. doi: 10.1103/PhysRevLett.77.3865.
- Sergey Pozdnyakov and Michele Ceriotti. Smooth, exact rotational symmetrization for deep learning on point clouds. 36:79469–79501. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/fb4a7e3522363907b26a86cc5be627ac-Abstract-Conference.html.
- E. Prodan and W. Kohn. Nearsightedness of electronic matter. 102(33):11635–11638. doi: 10.1073/pnas.0505436102.
- Benjamin Rhodes, Sander Vandenhaute, Vaidotas Šimkus, James Gin, Jonathan Godwin, Tim Duignan, and Mark Neumann. Orb-v3: Atomistic simulation at scale.

- Kristof Schütt, Pieter-Jan Kindermans, Huziel Enoc Saucedo Felix, Stefan Chmiela, Alexandre Tkatchenko, and Klaus-Robert Müller. SchNet: A continuous-filter convolutional neural network for modeling quantum interactions. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., a. URL https://papers.nips.cc/paper_files/paper/2017/hash/303ed4c69846ab36c2904d3ba8573050-Abstract.html.
- Kristof Schütt, Oliver Unke, and Michael Gastegger. Equivariant message passing for the prediction of tensorial properties and molecular spectra. In *Proceedings of the 38th International Conference on Machine Learning*, pp. 9377–9388. PMLR, b. URL <https://proceedings.mlr.press/v139/schutt21a.html>.
- Tess E. Smidt. Euclidean Symmetry and Equivariance in Machine Learning. 3(2):82–85. ISSN 2589-5974. doi: 10.1016/j.trechm.2020.10.006.
- Carsten G. Staacke, Hendrik H. Heenen, Christoph Scheurer, Gábor Csányi, Karsten Reuter, and Johannes T. Margraf. On the Role of Long-Range Electrostatics in Machine-Learned Interatomic Potentials for Complex Battery Materials. 4(11):12562–12569. doi: 10.1021/acsaem.1c02363.
- J. Tersoff. New empirical approach for the structure and energy of covalent systems. 37(12): 6991–7000. doi: 10.1103/PhysRevB.37.6991.
- Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation- and translation-equivariant neural networks for 3D point clouds. URL <https://arxiv.org/abs/1802.08219v3>.
- Alexandre Tkatchenko and Matthias Scheffler. Accurate Molecular Van Der Waals Interactions from Ground-State Electron Density and Free-Atom Reference Data. 102(7):073005. ISSN 0031-9007, 1079-7114. doi: 10.1103/physrevlett.102.073005.
- Oliver T. Unke and Hartmut Maennel. E3x: E(3)-Equivariant Deep Learning Made Easy.
- Oliver T. Unke and Markus Meuwly. PhysNet: A Neural Network for Predicting Energies, Forces, Dipole Moments, and Partial Charges. 15(6):3678–3693. ISSN 1549-9618. doi: 10.1021/acs.jctc.9b00181.
- Oliver T. Unke, Stefan Chmiela, Michael Gastegger, Kristof T. Schütt, Huziel E. Saucedo, and Klaus-Robert Müller. SpookyNet: Learning force fields with electronic degrees of freedom and nonlocal effects. 12(1), a. ISSN 2041-1723. doi: 10.1038/s41467-021-27504-0.
- Oliver T. Unke, Stefan Chmiela, Huziel E. Saucedo, Michael Gastegger, Igor Poltavsky, Kristof T. Schütt, Alexandre Tkatchenko, and Klaus-Robert Müller. Machine Learning Force Fields. 121 (16):10142–10186, b. ISSN 0009-2665. doi: 10.1021/acs.chemrev.0c01111.
- Brandon M. Wood, Misko Dzamba, Xiang Fu, Meng Gao, Muhammed Shuaibi, Luis Barroso-Luque, Kareem Abdelmaqsoud, Vahe Gharakhanyan, John R. Kitchin, Daniel S. Levine, Kyle Michel, Anuroop Sriram, Taco Cohen, Abhishek Das, Ammar Rizvi, Sushree Jagriti Sahoo, Zachary W. Ulissi, and C. Lawrence Zitnick. UMA: A Family of Universal Models for Atoms.
- Tian Xie and Jeffrey C. Grossman. Crystal Graph Convolutional Neural Networks for an Accurate and Interpretable Prediction of Material Properties. 120(14):145301. doi: 10.1103/PhysRevLett.120.145301.
- Yang You, Jing Li, Sashank Reddi, Jonathan Hseu, Sanjiv Kumar, Srinadh Bhojanapalli, Xiaodan Song, James Demmel, Kurt Keutzer, and Cho-Jui Hsieh. Large Batch Optimization for Deep Learning: Training BERT in 76 minutes. URL <https://openreview.net/forum?id=Syx4wnEtVH>.
- Maksim Zhdanov, Max Welling, and Jan-Willem van de Meent. Erwin: A Tree-based Hierarchical Transformer for Large-scale Physical Systems. URL <https://openreview.net/forum?id=MrphqqwnKv>.

A INVARIANCE OF LOREM

As explained in Section 4, the potential at each atom, $\mathbf{V}_i^l$ is equivariant because the long-range message passing step consists of equivariant operations: Addition, and multiplication with a factor shared per l . In practical implementations of Ewald summation in periodic systems, there is one extra step: To prevent divergences, the total charge must be zero, which can be done either by simply subtracting the sum of the total charge from each charge, or equivalently, by subtracting an analytical correction from the potentials. Both approaches are equivalent, and since summation is equivariant, also do not break invariance. For this reason, the total procedure retains equivariance.

To numerically confirm these considerations, we compute the mean absolute error of energy predictions over a degree $L = 3$ Lebedev grid of rotations, plus inversions, with respect to the unrotated case, for the first 5 structures of the MgO surface validation set. The errors are in line with precision expectations: For single precision, $1.031 \cdot 10^{-6}$ meV and for double precision $5.913 \cdot 10^{-15}$ meV.

B DATASET CONSTRUCTION

MgO surface The MgO surface dataset from Ko et al. is used without modification. A representative snapshot of the unit cell is shown in Fig. 2A. The dataset is built from four configuration types, each derived from a distinct initial structure:

1. A pure MgO surface with the Au_2 dimer oriented perpendicular to the surface,
2. A pure MgO surface with the Au_2 dimer oriented parallel to the surface,
3. An Al-doped MgO surface with the Au_2 dimer perpendicular to the surface, and
4. An Al-doped MgO surface with the Au_2 dimer parallel to the surface.

Each of these initial configurations was first geometry optimized. For the two perpendicular, ‘non-wetting’, cases (1 and 3), the distance between the lower Au atom and the O atom directly beneath it was systematically varied, displacing the Au_2 dimer as a whole. From these distance-dependent samples, a subset was randomly selected. To introduce structural diversity, Gaussian noise was applied to each configuration: a standard deviation of 0.02 Å for atoms in the MgO substrate and 0.1 Å for the gold cluster. After perturbation, 1250 structures were selected from each of the four configuration types, yielding a total of 5000 samples. For these, energies and forces were computed using the FHI-aims code (Abbott et al.) and the Perdew–Burke–Ernzerhof (PBE) functional (Perdew et al.). A random 90/10 train–validation split was used, as reported in Ko et al.. Additional energy–distance curves with equal spacing were constructed for the perpendicular cases and are shown in Fig. 2B.

NaCl cluster The NaCl cluster dataset from Ko et al. is used without modification. First, the Na_9Cl_8^+ cluster was optimized in vacuum. A snapshot is shown in the top panel of Fig. 2C. From this geometry, the Na atom farthest from all other Na atoms was removed, yielding the Na_8Cl_8^+ cluster shown in the bottom panel of Fig. 2C. From these two structures, additional configurations were created by varying the distance between a selected pair of Na atoms along the line connecting them. In Fig. 2C, this moving atom is illustrated by transparent copies along its trajectory. Training datasets for both Na_9Cl_8^+ and Na_8Cl_8^+ were constructed by randomly sampling configurations along these trajectories. Gaussian noise with a standard deviation of 0.05 Å was applied to the atomic coordinates, resulting in 2500 perturbed structures for each molecule and a total of 5000 instances. Energies and forces were computed using the FHI-aims code and the PBE functional. A random 90/10 train–validation split was used, as reported by Ko et al.. Additional energy–distance curves with equal spacing were constructed and are shown in Fig. 2D.

Cumulenes The cumulene dataset from Unke et al. (b) is used without modification. The geometry of the linear molecule is shown in Fig. 3A. The dataset contains 4500 randomly sampled cumulene structures with nine carbon atoms, divided into training, validation, and test sets with 2000/500/2000 instances, respectively. The energy curves shown in Fig. 3B are based on a controlled subset in which only one terminal carbon atom is rotated, while the rest of the molecule remains fixed. A visualization of this rotational motion is provided in the inset of Fig. 3B.

Biodimers The Biodimers dataset from Huguenin-Dumittan et al.; Burns et al. (2017) is used without modification. It consists of 2291 relaxed organic sidechain–sidechain fragments, including small molecules such as ethanol, acetamide, and others. Based on molecular properties, the dataset is divided into six categories: each molecule is classified as either polar, apolar, or charged, resulting in the following dimer types—polar–polar (PP), polar–apolar (PA), polar–charged (PC), apolar–apolar (AA), apolar–charged (AC), and charged–charged (CC). A representative CC dimer is shown in Fig. 4A. The initial separation between molecules reflects their positions in protein structures. From this configuration, the intermolecular distance is incrementally increased up to 15 Å, resulting in a total of 29 783 dimer instances. Instances where the final separation exceeds the initial distance by more than 4 Å (13 743 samples) are designated as the test set. The remaining 16 040 instances form the training set. Energies and forces were computed using the FHI-aims code and the HSE06 hybrid functional. For each of the six dimer types, a random 80/20 train–validation split was applied. The resulting subsets were then merged into a single training set and a single validation set.

S_N2 reactions The S_N2 reactions dataset from Unke & Meuwly; Frank et al. (a) is used without modification. It contains molecular structures in vacuum that model nucleophilic substitution (S_N2) reactions. The dataset includes molecules of the types XCH₃Y[−], CH₃X, HX, CHX, CH₂X, XY, X, and Y, for all possible combinations of X, Y ∈ {F, Cl, Br, I}. A representative reaction coordinate with corresponding snapshots is shown in Fig. 5. Additional species such as H₂, CH₂, and CH₃ are also included in the dataset. Training configurations were generated using ab initio molecular dynamics simulations at 5000 K, with a time step of 0.1 fs. Full computational details, including the level of theory, are provided in the original publications. The dataset is randomly split into 405 000 training, 5000 validation, and 42 708 test samples.

C EWALD SUMMATION

The evaluation of inverse power law potentials has been the object of much study in computational physics and chemistry. The general task is to compute the potential V and the potential energy E :

$$V(\mathbf{r}) = \sum_{j=1}^N \sum_{\mathbf{n} \in \mathbb{Z}^3} \frac{q_j}{|\mathbf{r} - (\mathbf{r}_j + \mathbf{C}\mathbf{n})|^p} \quad E_{\text{LR}} = \sum_i \frac{1}{2} q_i V(\mathbf{r}_i), \quad (2)$$

where q_i are charges, or other coefficients, placed on each atomic position, and $\mathbf{n}$ is a vector indicating periodic replicas. For $p \leq 3$, this infinite sum converges conditionally, i.e., convergence or divergence depends on the summation order. Ewald summation was developed to tackle this problem for electrostatics ($p = 1$) and later extended to other exponents; its basic concept is to split $1/r$ into a short-ranged part, which converges fast in real space, and a long-range part, which is smooth, and therefore converges well, in reciprocal space.

A naive implementation of this method scales $O(N^2)$, which can be brought down to $O(N^{3/2})$ by choosing the cutoffs to be proportional to the size of the simulation cell. Particle–mesh Ewald (PME, P3M) algorithms reduce this to $O(N \log N)$ by interpolating charges onto a grid and employing the fast Fourier transform for the reciprocal-space part (Darden et al.; Hockney & Eastwood). While such methods are standard in force fields, implementations in popular machine learning frameworks, PyTorch (Paszke et al.) and JAX (Bradbury et al.), have only become available recently (Loche et al.). In non-periodic systems, the naive evaluation of Eq. (2) scales $O(N^2)$. However, methods like the fast multipole expansion can reduce this cost to $O(N)$ (Greengard & Rokhlin).

D ADDITIONAL RESULTS FOR CUMULENE

In Fig. 6, we present the curves related to Table 2. All models where the dihedral information is accessible can resolve the benchmark, but vary in match to the ground truth. In particular, we note that models with two short-range message passing steps perform better than a single one. This may be due to the high degree of degeneracy in atomic environments at short cutoffs, which is alleviated by message passing.




Figure 6: Rotational profile of cumulene for variations of the LOREM model with different r_c , different numbers of short-range message passing steps, and with and without long-range message passing. All curves that are identical to zero have been collapsed into a single one for readability.

Table 2: Ability of different LOREM models, with and without long-range message passing and with different numbers of short-range message passing steps, to solve the cumulene benchmark task. A tick (✓) indicates yes, a cross (×) indicates no.

Cutoff	1×SR + LR	2×SR + LR	1×SR	2×SR	3×SR	4×SR
2.5 Å	✓	✓	×	×	×	✓
3.0 Å	✓	✓	×	×	✓	✓
3.5 Å	✓	✓	×	✓	✓	✓

E LOREM BLOCKS DESCRIPTION

Below is the description of the blocks that were not described in the main text:

Spherical norm For any set of spherical (equivariant) features, i.e., possessing an index l and for each l a corresponding index m (see Section 2), the spherical norm is defined as:

$$\|\mathbf{S}^l\|_2 = \sqrt{(2l+1)^{1/2} \sum_{m=-l}^l S_{lm}^2}.$$

We find that the empirically determined prefactor reduces variance across l .

Update block Updates $\mathbf{Y}$ to scalar features $\mathbf{X}$ are processed with an update block:

$$\begin{aligned}\mathbf{X} &\leftarrow \mathbf{X} + \text{MLP}(\mathbf{Y}) \\ \mathbf{X} &\leftarrow \text{LayerNorm}(\mathbf{X}) \\ \mathbf{X} &\leftarrow \mathbf{X} + \text{MLP}(\mathbf{X}) \\ \mathbf{X} &\leftarrow \text{LayerNorm}(\mathbf{X}),\end{aligned}$$

following a residual structure.

F MODEL DESCRIPTION AND TRAINING SETUP

LOREM For all experiments, except the ones where variations were tested, the same LOREM model architecture (Section 5) was used: A cutoff radius of 5 Å, $l_{\max} = 6$ for spherical features and $l_{\max} = 2$ for the long-range message passing, 128 scalar features, 8 channels for spherical features, and 32 radial basis functions. We perform only the initial short-range message passing step. This model has 1 021 198 learnable parameters. It was implemented using the `e3x` library (Unke & Maennel). Training is performed entirely in `float32` precision; while we find that reduced precision has only a minor effect on training dynamics, it can significantly alter validation and test results.

Training parameters vary slightly between datasets and were chosen by hand to minimize validation error;⁴ the exact parameters can be found in Table 3. We report errors for the better out of two runs with different seeds; the presented conclusions hold for both models. In all cases, except NaCl⁵, the learning rate was decayed linearly after 10 epochs, from the indicated starting value to $1 \cdot 10^{-6}$; the ADAM (Kingma & Ba) and LAMB (You et al.) optimizers were used. The loss function was a simple mean squared error, with the energy residuals normalized by number of atoms. The squared residuals were averaged over the whole batch, including over atoms and components in the case of forces, and then summed and weighted with a factor. The checkpoint with the lowest summed R^2 of energy and forces, evaluated on the validation set, was used for experiments. Training times are given for the entire run, not the time until the best checkpoint.

⁴We note that the validation error in the cumulene dataset does not readily correlate with ability to predict the dihedral curve in Fig. 3 – while all models are able to resolve the task, the agreement varies.

⁵Here, exponential learning rate decay was employed.

Dataset	Optimizer	Initial LR	Epochs	Batch size	E weight	F weight	Time
MgO surface	ADAM	$4 \cdot 10^{-4}$	4000	32	1000	1.0	17 h
Biodimers	ADAM	$1 \cdot 10^{-4}$	4000	32	0.5	0.5	42 h
Cumulene	LAMB	$1 \cdot 10^{-3}$	2000	32	0.5	0.5	2 h
NaCl cluster	ADAM	$1 \cdot 10^{-3}$	8000	64	0.5	0.5	7 h
$\text{S}_{\text{N}}2$ reactions	ADAM	$1 \cdot 10^{-4}$	500	32	0.5	0.5	17 h

Table 3: Training settings for LOREM for different datasets. Training times are given for a single Nvidia H100 SXM5 GPU.

MACE The MACE MLIP (Batatia et al. (b)) is an equivariant message-passing neural network as described in Section 2. It is used with the standard setting of two message-passing steps, i.e., an effective cutoff of $10 \text{ \AA}$. Training generally used default hyperparameters, with some adjustments made on hyperparameters related to training dynamics. Specifically, the cumulene system was trained using the default energy-to-forces loss weight ratio of 1:100 and ‘128x0e + 128x1o + 128x2e’ hidden irreps. All other systems were trained using a 1:10 ratio and ‘128x0e + 128x1o’ hidden irreps.

All experiments, except those involving cumulenes and the NaCl cluster, used the SWA protocol, which swaps the loss weights between energy and forces at a specified epoch. This epoch was set where the energy RMSE plateaued, with the subsequent training continued until loss saturation. All models were trained with $l_{\text{max}} = 2$, a batch size of 32, and utilized cuEquivariance acceleration.

PET PET (Pozdnyakov & Ceriotti) is an unconstrained transformer model, consisting of multiple edge-to-edge transformer layers within local neighborhoods followed by message passing. Similar to MACE, it is also used with two message-passing steps and an effective cutoff of $10 \text{ \AA}$. To ensure approximate rotational invariance, it is trained with data augmentation. No inference-time symmetrization is used in our experiments.

Unless otherwise specified, we used default hyperparameters (cutoff radius of $5 \text{ \AA}$ with cosine smoothing over the outermost $0.2 \text{ \AA}$, $d_{\text{PET}} = 128$, $d_{\text{head}} = 128$, $d_{\text{feedforward}} = 512$, with 8 heads per attention layer, and 2 attention layers per GNN layer), with some adjustments made related to training dynamics. Models were trained using an epoch-based scheduler, which halved the learning rate after 250 epochs. This applies to all datasets except biodimers. For biodimers, a ReduceLROnPlateau scheduler was used instead, reducing the learning rate by 20% if the loss did not improve for 100 consecutive epochs. To improve training stability, biodimers training also employed gradient clipping with a maximum gradient norm of 5. Every training run included 10 warmup epochs, during which the learning rate was linearly increased from zero to the preset learning rate of 1×10^{-4} .

For the biodimers and NaCl cluster datasets, the targets were normalized by their standard deviation in the training set. The energy-to-forces loss weight ratio was set to 1:1 for biodimers and MgO surface datasets, while the NaCl cluster dataset used a ratio of 1:10.

For cumulene, PET was trained for 20 100 epochs with a batch size of 16, equal energy and forces weight, and a maximum learning rate of $2 \cdot 10^{-4}$. The learning rate was increased linearly from 0 over 100 epochs, and then reduced by 10% every 500 epochs over the entire training run. Architecturally, 5 attention layers and cosine cutoff function was employed to resolve the cumulenes. We found that increasing the number of attention layers, and adapting the cutoff function, were critical to resolve this benchmark task.

CACE-LES The CACE-LES model was originally presented in Cheng (a), with an additional modification introducing a long-range component in Cheng (b). Similar to MACE, CACE-LES is an equivariant message-passing neural network. After short-range message passing, invariant scalar pseudo-charges are predicted and passed to the long-range part of Ewald summation; the resulting energy contribution is added to the energy predictions of the short-range message-passing model. We use this model with the recommended setting of one message-passing step and, depending on the system, the following corresponding effective cutoff radii: $5.5 \text{ \AA}$ for the MgO surface, $5.29 \text{ \AA}$ for the NaCl cluster—these are the settings used in the original paper Kim et al.—and we chose a cutoff of $5 \text{ \AA}$ for biodimers and cumulene.

For the MgO surface and NaCl cluster datasets, models were taken from Kim et al.. For the cumulenes and biodimers datasets, the hyperparameters were as follows: 6 Bessel radial functions, $c = 8$, $l_{\max} = 3$, $\nu_{\max}$, $N_{\text{embedding}} = 2$, one message-passing layer, one-dimensional hidden variable, $\sigma = 1$, and $dl = 2$. Training followed the example in the original paper: the first 200 epochs used an energy loss weight of 0.1 and a forces loss weight of 1000, after which the energy loss weight was changed to 1, 10, and 1000 every 100 epochs, yielding a total of 500 epochs. The learning rate was 5×10^{-3} , with a step learning rate schedule decreasing it by a factor of 2 every 20 steps.

4G-NN The 4G-NN model was designed to tackle the benchmarks introduced by Ko et al., and introduced in that work. It consists of a shallow neural network acting on rotationally invariant features, predicting both a local energy contribution and an electronegativity, which is then used in a charge equilibration procedure that globally redistributes charges to minimize an energy expression. This process can be seen as a physics-inspired long-range message passing scheme iterated until a fixed point is reached. As the model requires charge labels to train, which are not available for all datasets, we did not train this model for our experiments but instead include results from Ko et al. Results for the 4G-NN model are only available for the NaCl cluster and MgO surface datasets. The cutoffs used in the original paper are as follows: 4.23 Å for the MgO surface and 5.29 Å for the NaCl cluster.