

---

# Equivariant Message Passing for the Prediction of Tensorial Properties and Molecular Spectra

---

Kristof T. Schütt<sup>1,2</sup> Oliver T. Unke<sup>1,2</sup> Michael Gastegger<sup>1,3</sup>

## Abstract

Message passing neural networks have become a method of choice for learning on graphs, in particular the prediction of chemical properties and the acceleration of molecular dynamics studies. While they readily scale to large training data sets, previous approaches have proven to be less data efficient than kernel methods. We identify limitations of invariant representations as a major reason and extend the message passing formulation to rotationally equivariant representations. On this basis, we propose the *polarizable atom interaction neural network* (PAINN) and improve on common molecule benchmarks over previous networks, while reducing model size and inference time. We leverage the equivariant atomwise representations obtained by PAINN for the prediction of tensorial properties. Finally, we apply this to the simulation of molecular spectra, achieving speedups of 4-5 orders of magnitude compared to the electronic structure reference.

## 1. Introduction

Studying dynamics of chemical systems allows insight into processes such as reactions or the folding of proteins, and constitutes a fundamental challenge in computational chemistry. Since the motion of atoms is governed by the laws of quantum mechanics, accurate *ab initio* molecular dynamics (MD) simulations may require solving the Schrödinger equation for millions of time steps. While the exact solution is infeasible to compute for all but the smallest systems, even fast approximations such as density functional the-

ory quickly become prohibitive for large systems and the prediction of accurate spectra.

Recently, machine learning potentials (Behler, 2016; Unke et al., 2020; von Lilienfeld et al., 2020) have gained popularity for studying systems ranging from small molecules at high levels of theory (Chmiela et al., 2018; Westermayr et al., 2020) to systems with thousands or millions of atoms (Morawietz et al., 2016; Bartók et al., 2018; Lu et al., 2020). In particular, message-passing neural networks (Gilmer et al., 2017) (MPNNs) yield accurate predictions for chemical properties across chemical compound space and can handle large amounts of training data. Albeit MPNNs have significantly increased in accuracy over the years (as well as in computational cost), kernel methods with manually crafted features (Chmiela et al., 2017; Christensen et al., 2020; Bartók et al., 2010) have still proven to perform better when only small training sets are available.

While molecules are often represented as graphs, they are in fact interacting particles in a continuous 3d space. Consequently, SchNet (Schütt et al., 2017) modeled message passes as continuous-filter convolutions over that space, albeit with rotationally invariant filters. As Miller et al. (2020) pointed out, this leads to a loss of relevant directional, equivariant information. Klicpera et al. (2020a) have introduced directional message-passing, the angular information here is restricted to the messages while the representation of nodes (atoms) remains rotationally invariant. While equivariant convolutions have been successfully applied in computer vision (Cohen & Welling, 2017; Weiler et al., 2018b; Worrall & Brostow, 2018), previous approaches to molecular prediction (Thomas et al., 2018; Anderson et al., 2019) have not reached the accuracy of their rotationally invariant counterparts.

In this work, we propose rotationally equivariant message passing and the *polarizable atom interaction neural network* (PAINN) architecture as one instance of it. We examine the limited capability of rotation-invariant representations to propagate directional information and show that equivariant representations do not suffer from this issue. PAINN outperforms invariant message passing networks on common molecular benchmarks and performs at small sample sizes on par with kernel methods that have been deemed to be

---

<sup>1</sup>Machine Learning Group, Technische Universität Berlin, 10587 Berlin, Germany <sup>2</sup>Berlin Institute for the Foundations of Learning and Data, 10587 Berlin, Germany <sup>3</sup>BASLEARN – TU Berlin / BASF Joint Lab for Machine Learning, 10587 Berlin, Germany. Correspondence to: Kristof T. Schütt <kristof.schuett@tu-berlin.de>, Michael Gastegger <michael.gastegger@tu-berlin.de>.

more data-efficient than neural networks. Beyond that, the rotationally equivariant representation of PAiNN enables the prediction of tensorial properties which we apply to the ring-polymer MD simulation of infrared and Raman spectra. By an acceleration of 4-5 orders of magnitude, PAiNN makes the simulation of these spectra feasible, reducing the runtime in one case from projected 25 years to one hour.

## 2. Related work

Behler & Parrinello (2007) introduced neural network potentials taking atom-centered symmetry functions based on distances and angles as features. Graph neural networks (Scarselli et al., 2008) for molecular graphs (Duvinaud et al., 2015; Kearnes et al., 2016) and 3d geometries (Schütt et al., 2017; Gilmer et al., 2017; Schütt et al., 2018; Unke & Meuwly, 2019; Lubbers et al., 2018) do not require such manually crafted features, but learn embeddings of atom types and use graph convolutions or, more general, message passing (Gilmer et al., 2017), to model atom interactions based on interatomic distances. Klicpera et al. (2020a) introduced directional message passing by including additional angular information in the message.

Steerable CNNs (Cohen & Welling, 2017; Weiler et al., 2018a) allow to build equivariant filter banks according to known symmetries of associated feature types. While these approaches work on grids, architectures such as Tensor Field Networks (Thomas et al., 2018), Cormorant (Anderson et al., 2019) and NEQUIP (Batzner et al., 2021) use equivariant convolutions based on spherical harmonics (SH) and Clebsch-Gordon (CG) transforms for point clouds. Kondor & Trivedi (2018) have described a general framework on equivariance and convolutions in neural networks focusing on irreducible representations. In contrast, PAiNN models equivariant interactions in Cartesian space which is conceptually simpler and does not require tensor contractions with CG coefficients. A similar approach was proposed by Jing et al. (2021) with the GVP-GNN. Both approaches are designed for distinct applications resulting in crucial differences in the design of the neural network architectures, most notably the message functions. The GVP-GNN has been designed for single-point protein sequence predictions, e.g., allowing the use of non-smooth components. In contrast, PAiNN is designed for simulations with millions of inference steps, requiring a fast message function and a smoothly differentiable model.

## 3. Equivariant message passing

**Notation.** To clearly distinguish between the two concepts, we write feature vectors as  $\mathbf{x} \in \mathbb{R}^{F \times 1}$  and vectors in 3d coordinate space as  $\vec{r} \in \mathbb{R}^{1 \times 3}$ . Vectorial features will be written as  $\vec{\mathbf{x}} \in \mathbb{R}^{F \times 3}$ . Norms  $\|\cdot\|_k$ , scalar products  $\langle \cdot, \cdot \rangle$  and

tensor products are calculated along the spatial dimension. All other operations are calculated along the feature dimension, if not stated otherwise. We write the Hadamard product as  $\odot$ .

### 3.1. Message passing for 3d-embedded graphs

MPNNs build complex representations of nodes within their local neighborhood in a graph through a repeated exchange of *messages* followed by *updates* of node features. Here, we consider graphs embedded in 3d Euclidean space, where edges are specified by the relative positions  $\vec{r}_{ij} = \vec{r}_j - \vec{r}_i$  of nodes  $i, j$  within a local neighborhood  $N(i) = \{j \mid \|\vec{r}_{ij}\| < r_{\text{cut}}\}$ . In a chemistry context, this agrees with the fact that a large part of the energy variations can be attributed to local interactions, often conceptualized as bond lengths and angles. Thus, only nodes within that range can interact directly, so the number of messages does not scale quadratically with the number of the nodes. A general MPNN for embedded graphs can be written as

$$\mathbf{m}_i^{t+1} = \sum_{j \in N(i)} \mathbf{M}_t(\mathbf{s}_i^t, \mathbf{s}_j^t, \vec{r}_{ij}) \quad (1)$$

$$\mathbf{s}_i^{t+1} = \mathbf{U}_t(\mathbf{s}_i^t, \mathbf{m}_i^{t+1}). \quad (2)$$

with the update and message functions  $\mathbf{U}_t$  and  $\mathbf{M}_t$ , respectively (Gilmer et al., 2017). Many neural network potentials and even conventional CNNs can be cast in this general framework. Rotational invariance of the representation can be ensured by choosing rotationally invariant message and update functions, which by definition need to fulfill

$$\mathbf{f}(\vec{\mathbf{x}}) = \mathbf{f}(R\vec{\mathbf{x}}), \quad (3)$$

for any rotation matrix  $R \in \mathbb{R}^{3 \times 3}$ .

### 3.2. Building equivariant MPNNs

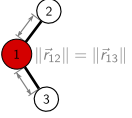
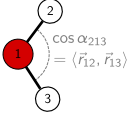
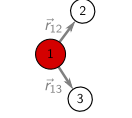
To obtain more expressive representations of local environments, neurons do not have to be scalar, but can be geometric objects such as vectors and tensors (Hinton et al., 2011; Cohen & Welling, 2017; Thomas et al., 2018; Anderson et al., 2019). For the purpose of this work, we restrict ourselves to scalar and vectorial representations  $\mathbf{s}_i^t$  and  $\vec{\mathbf{v}}_i^t$ , respectively, such that a corresponding message pass can be written as

$$\vec{\mathbf{m}}_i^{v;t+1} = \sum_{j \in N(i)} \vec{\mathbf{M}}_t(\mathbf{s}_i^t, \mathbf{s}_j^t, \vec{\mathbf{v}}_i^t, \vec{\mathbf{v}}_j^t, \vec{r}_{ij}) \quad (4)$$

Message and update functions for scalar features can be defined analogously. Rotational equivariance of vector features  $\vec{\mathbf{v}}_i^{t+1}$  can be ensured by employing rotationally equivariant functions  $\vec{\mathbf{U}}_t$  and  $\vec{\mathbf{M}}_t$ , fulfilling

$$R\vec{\mathbf{f}}(\vec{\mathbf{x}}) = \vec{\mathbf{f}}(R\vec{\mathbf{x}}) \quad (5)$$

Table 1. Comparison of expressiveness and computational complexity of distances, angles and directions for a simple message function of the oxygen atom of a water molecule.

Features	Distances	Angles	Directions
$H_2O$			
Message $M$ at atom $i$	$\sum_{j \in \mathcal{N}_i} k_{F_{ij}} k$	$\sum_{j \in \mathcal{N}_i} \sum_{k \in \mathcal{N}_i} j i k$	$\sum_{j \in \mathcal{N}_i} \frac{F_{ij}}{k_{F_{ij}} k}$
Scaling with neighbors	$O(jNj)$	$O(jNj^2)$	$O(jNj)$
Resolve change of $k_{F_{ij}} k$	yes	no	no
Resolve change of $\alpha_{213}$	no	yes	yes

for any rotation matrix  $R \in \mathbb{R}^{3 \times 3}$ , where the matrix-vector product is applied over the spatial dimension. This constitutes essentially a linearity constraint for directional information. Therefore, any nonlinearities of the model need to be applied to scalar features. In particular, equivariant MPNNs may use the following operations:

- Any (nonlinear) function of scalars:  $f(s)$
- Scaling of vectors:  $s \vec{v}$
- Linear combinations of equivariant vectors:  $\mathbf{W}\vec{v}$
- Scalar products:  $s = k\vec{v}k^2$ ,  $s = h\vec{v}_1, \vec{v}_2$
- Vector products:  $\vec{v}_1 \times \vec{v}_2$

Thus, directional information is preserved during message passing while invariant representations can be recovered from scalar products to obtain rotationally invariant predictions.

### 3.3. Limits of rotationally invariant representations

In the following, we examine the expressiveness of invariant and equivariant representations at the example of simple 2d molecular structures. To ensure rotational invariance of a predicted property, the message function is often restricted to depend only on rotationally invariant inputs such as distances (Schütt et al., 2017; Lubbers et al., 2018) or angles (Klicpera et al., 2020a;b). While a single pass is limited to interactions within a local environment, successive message passes are supposed to construct more complex representations and propagate local information beyond the neighborhood. This raises the question whether rotationally invariant representations of atomic environments  $\mathbf{h}_i$  are sufficient here.




Figure 1. Illustration of message passing using angles and directions for two structures. All edges within the cutoff range (dashed lines) have equal length. The representations of the blue and red node are the same using angles (left), while directions allow to distinguish both structures (right).

Tab. 1 compares three simplified message functions for the example of a water molecule. Using a distance-based message function, the representation of the atom 1 (oxygen) is able to resolve changing bond lengths to atoms 2 and 3 (hydrogens), however it is not sensitive to changes of the bond angle. On the other hand, using the angle directly as part of the message function can not resolve the distances. Therefore, a combination of distances and angles is required to obtain a more expressive message function (Klicpera et al., 2020a). Unfortunately, including angles in the messages scales  $O(jNj^2)$  with the number of neighbors. Alternatively, directions to neighboring atoms may be used as messages  $\vec{M}_t(\vec{r}_{ij}) = \vec{r}_{ij} / k\vec{r}_{ij}k$ . Employing the update function  $U_t(m) = kmk^2$ , this is related to angles as follows:

$$\left\| \sum_{j=1}^N \frac{\vec{r}_{ij}}{k\vec{r}_{ij}k} \right\|^2 = \sum_{j,k} \left\langle \frac{\vec{r}_{ij}}{k\vec{r}_{ij}k}, \frac{\vec{r}_{ik}}{k\vec{r}_{ik}k} \right\rangle = \sum_{j=1}^N \sum_{k=1}^N \cos \alpha_{jik}.$$

Thus, using equivariant messages, the runtime complexity remains  $O(jNj)$  while angular information can be resolved. Note that this update function contracts the equivariant messages to a rotationally invariant representation.

Beyond the computational benefits, equivariant representations allow to propagate directional information beyond the neighborhood which is *not* possible in the invariant case. Fig. 1 illustrates this at a minimal example, where four atoms are arranged in an equidistant chain with the cutoff chosen such that only neighboring atoms are within range. We observe that for the two arrangements the angles are equal as well (Fig. 1, left). Therefore, they are indistinguishable for invariant message passing with distances and angles. In contrast, the equivariant representations differ in sign of their components (Fig. 1, right). When contracting them to invariant representations, as in the previous example, this information is lost. However, we may instead retain equivariance in the representation and design a message function  $\vec{M}_t$  that does not only propagate directions to neighboring atoms, but those of equivariant representations as well. This enables the efficient propagation of directional information by scaling linearly with the number of neighbors and keeping the required cutoff small.

Note that neither many-body representations, using angles, dihedral angles etc., nor equivariant representations corresponding to a multipole expansion are complete. It has been shown that even when including up to 4-body invariants, there are structures with as few as eight atoms which cannot be distinguished (Pozdnyakov et al., 2020) and that a many-body expansion up to  $n$ -body contributions is necessary to guarantee convergence for a system consisting of  $n$  atoms (Hermann et al., 2007). The same holds for multipole expansions, where scalars and vectors correspond to the 0th and 1st order. Thus, even spherical harmonics expansions with less than infinite degree are unable to represent arbitrary equivariant  $n$ -body functions. Instead, a practically sufficient and computationally efficient approach for the problem at hand is desirable.

In the following, we propose a neural network architecture that takes advantage of these properties. It overcomes the limitations of invariant representations discussed above, which we will demonstrate at the example of an organometallic compound in Section 5.3.2.

#### 4. Polarizable atom interaction neural network (PAINN)

The potential energy surface  $E(Z_1, \dots, Z_N, \vec{r}_1, \dots, \vec{r}_N)$ , with nuclear charges  $Z_i \in \mathbb{N}$  and atom positions  $\vec{r}_i \in \mathbb{R}^3$  exhibits certain symmetries. Generally, they include invariance of the energy towards the permutation of atom indices  $i$ , as well as rotations and translations of the molecule. A neural network potential should encode these constraints to ensure the symmetries of the predicted energy surface and increase data efficiency. A common inductive bias of the neural network potential is a decomposition of the energy into atomwise contributions  $E = \sum_{i=1}^N \epsilon(\mathbf{s}_i)$ , where an output network  $\epsilon$  predicts energy contributions from atoms embedded within their chemical environment, represented by  $\mathbf{s}_i \in \mathbb{R}^F$  (Behler & Parrinello, 2007; Bartók et al., 2013).

While properties of chemical compounds may be such rotationally invariant scalars, they can also be equivariant tensorial properties, e.g. the multipole expansion of the electron density

$$n(\vec{r}) = q + \vec{\mu}^T \vec{r} + \vec{r}^T Q \vec{r} + \dots, \quad (6)$$

with charge  $q$ , dipole  $\vec{\mu}$ , quadrupole  $Q$  and so on. Similarly, one can interpret invariant and equivariant atomwise representations as local charges and dipole moments of atom-centered multipole expansions (Gastegger et al., 2020). We leverage this in Section 4.2 to predict tensorial molecular properties. We coin our proposed architecture *polarizable atom interaction neural network* (PAINN).

#### 4.1. Representation

The inputs to PAINN are the nuclear charges  $Z_i \in \mathbb{N}$  and positions  $\vec{r}_i \in \mathbb{R}^3$  for each atom  $i$ . Similar to previous approaches, the invariant atom representations are initialized to learned embeddings of the atom type  $\mathbf{s}_i^0 = \mathbf{a}_{Z_i} \in \mathbb{R}^{F-1}$ . We keep the number of features  $F$  constant throughout the network. The equivariant representations are set to  $\vec{\mathbf{v}}_i^0 = \vec{\mathbf{0}} \in \mathbb{R}^{F-3}$ , since there is no directional information available initially.

Next, we define message and update functions as introduced in Sec. 3. We use a residual structure of interchanging message and update blocks (Fig. 2a), resulting in coupled scalar and vectorial representations. For the residual of the scalar message function, we adopt the feature-wise, continuous-filter convolutions introduced by Schütt et al. (2017)

$$\begin{aligned} \mathbf{s}_i^m &= (\phi_s(\mathbf{s}) \quad W_s)_i \\ &= \sum_j \phi_s(\mathbf{s}_j) \quad W_s(k\vec{r}_{ij}k), \end{aligned} \quad (7)$$

where  $\phi_s$  consists of atomwise layers as shown in Fig. 2b. The rotationally-invariant filters  $W_s$  are linear combinations of radial basis functions  $\sin(\frac{r}{r_{\text{cut}}} k\vec{r}_{ij}k) / k\vec{r}_{ij}k$  as proposed by Klicpera et al. (2020a) with  $1 \leq n \leq 20$ . Additionally, we apply a cosine cutoff to the filters (Behler, 2011).

Analogously, we use continuous-filter convolutions for the residual of the equivariant message function

$$\begin{aligned} \vec{\mathbf{v}}_i^m &= \sum_j \vec{\mathbf{v}}_j \quad W_{vv}(\mathbf{s}_j) \quad W_{vv}(k\vec{r}_{ij}k) \\ &+ \sum_j v_s(\mathbf{s}_j) \quad W_{vs}^0(k\vec{r}_{ij}k) \frac{\vec{r}_{ij}}{k\vec{r}_{ij}k}, \end{aligned} \quad (8)$$

where the first term is a convolution of an invariant filter with scaled, equivariant features  $\sum_j \vec{\mathbf{v}}_j \quad v_s(\mathbf{s}_j)$ , related to the gating proposed by Weiler et al. (2018a) as an equivariant nonlinearity. This propagates directional information obtained in previous message passes to neighboring atoms. The second term is a convolution of invariant features with an equivariant filter. This can be derived as the gradient of an invariant filter

$$\frac{\partial}{\partial \vec{r}} W_{vs}(k\vec{r}_{ij}k) = W_{vs}^0(k\vec{r}_{ij}k) \frac{\vec{r}_{ij}}{k\vec{r}_{ij}k}.$$

Since  $W_{vs}^0(k\vec{r}_{ij}k)$  is another invariant filter, we can model it directly without taking the derivative. Furthermore, we use a shared network  $\phi$  to perform the transform of all message functions and split the features afterwards (see Fig. 2b).

After the features-wise message blocks, the update blocks are applied atomwise across features. The residual of the scalar update function is given by

$$\begin{aligned} \mathbf{s}_i^u &= \mathbf{a}_{ss}(\mathbf{s}_i, k\mathbf{V}\vec{\mathbf{v}}_i k) \\ &+ \mathbf{a}_{sv}(\mathbf{s}_i, k\mathbf{V}\vec{\mathbf{v}}_i k) \hbar \mathbf{U}\vec{\mathbf{v}}_i, \mathbf{V}\vec{\mathbf{v}}_i. \end{aligned} \quad (9)$$




Figure 2. The architecture of PAINN with the full architecture (a) as well as the message (b) and update blocks (c) of the equivariant message passing. In all experiments, we use 128 features for  $s_i$  and  $\vec{v}_i$  throughout the architecture. Other layer sizes are annotated in grey.




Figure 3. Gated equivariant block.

Again, we use a scaling functions computed by a shared network  $\mathbf{a}(s_i, k\mathbf{V}\vec{v}_i k)$  as nonlinearity. In this case, the norm of a linear combination of features is also used to obtain the scaling, thus, coupling the scalar representations with contracted equivariant features. In a second term, we use the scalar product of two linear combinations of equivariant features. Similarly, we define the residual for the equivariant features

$$\vec{v}_i^u = \mathbf{a}_{vv}(s_i, k\mathbf{V}\vec{v}_i k) \mathbf{U}\vec{v}_i, \quad (10)$$

which is again a nonlinear scaling of linearly combined equivariant features.

## 4.2. Prediction of tensorial properties

PAINN is able to predict scalar properties using atomwise output layers of rotationally invariant representations  $s_i$ , in the same manner as previous approaches (Schütt et al., 2017; Unke & Meuwly, 2019). Beyond that, the additional equivariant representations enable the prediction of equivariant

tensors of order  $M$  that can be constructed using a rank-1 tensor decomposition

$$T = \sum_{i=1}^N \sum_{k=1}^R \lambda(s_i) \vec{v}(s_i)_{k:1} \vec{v}(s_i)_{k:M}. \quad (11)$$

The scalar and vectorial components of this decomposition can be obtained from an output network of *gated equivariant blocks*, as shown in Fig. 3. These components again make use gated equivariant nonlinearities (Weiler et al., 2018b) and show similarities to geometric vector perceptrons (GVP) (Jing et al., 2021). However, while GVP blocks are used as message functions, PAINN keeps the pairwise message functions comparatively lightweight. Instead complex transformations are restricted to the atomwise update function and tensor output network.

Depending on the chemical property of interest, one may want to include additional constraints or replace  $\vec{v}_{k:m}$  with the atom positions to include the global molecular geometry. In the following, we demonstrate how this can be employed for two tensorial properties that are important for the simulation of molecular spectra.

The molecular dipole moment  $\vec{\mu}$  is the response of the molecular energy to an electric field  $\mathcal{r}_{\mathcal{F}} E$  and, at the same time, the first moment of the electron density (see Eq. 6). It is often predicted using latent atomic charges (Gastegger et al., 2017):

$$\vec{\mu} = \sum_{i=1}^N q_{\text{atom}}(s_i) \vec{r}_i. \quad (12)$$

This assumes the center of mass at  $\vec{r} = \vec{0}$  for brevity. While this only leverages invariant representations, we can take

advantage of the equivariant features of PAINN. In this picture, the molecular dipole is constructed from polarized atoms, i.e. both atomic dipoles and charges (Veit et al., 2020), yielding

$$\vec{\mu} = \sum_{i=1}^N \vec{\mu}_{\text{atom}}(\vec{v}_i) + q_{\text{atom}}(\mathbf{s}_i) \vec{r}_i. \quad (13)$$

In terms of Eq. 11, this is a simple case with  $M = 1$ ,  $R = 2$ , where the first term predicts the latent local dipoles and the second term computes the dipole generated by the local charges.

As an example with order  $M = 2$ , we consider the polarizability tensor, which describes the response of the molecular dipole to an electric field  $J_{\mathcal{F}}(\vec{\mu}) = H_{\mathcal{F}}(E)$ . We construct polarizability tensors using

$$= \sum_{i=1}^N \alpha_0(\mathbf{s}_i) I_3 + \vec{v}(\vec{v}_i) \quad \vec{r}_i + \vec{r}_i \quad \vec{v}(\vec{v}_i), \quad (14)$$

where the first term of the sum models isotropic, atomwise polarizabilities. The other two terms add the anisotropic components of the polarizability tensor. Similar to the charges in Eq. 13, here the atom positions are used to incorporate the global structure of the molecule.

## 5. Results

PAINN has been implemented using PYTORCH (Paszke et al., 2019) and SCHNETPACK (Schütt et al., 2018). All models use  $F = 128$  and two output layers as shown in Fig. 2 (blue) with the type of output layers depending on the property of interest. They were trained using the Adam optimizer (Kingma & Ba, 2014), the squared loss and weight decay  $\lambda = 0.01$  (Loshchilov & Hutter, 2019), if not stated otherwise. We decay the learning rate by a factor of 0.5 if the validation loss plateaus. We apply exponential smoothing with factor 0.9 to the validation loss to reduce the impact of fluctuations which are particularly common when training with both energies and forces. Please refer to the supplement for further details on training parameters.

### 5.1. Chemical compound space

We use the QM9 dataset of 130k small organic molecules (Ramakrishnan et al., 2014) with up to nine heavy atoms to evaluate the performance of PAINN for the prediction of scalar properties across chemical compound space. We predict the magnitude of the dipole moment using Eq. 13 and the electronic spatial extent by

$$\langle R^2 \rangle = \sum_{i=1}^N q_{\text{atom}}(\mathbf{s}_i) k \vec{r}_i k^2,$$

as implemented by SCHNETPACK. The remaining properties are predicted as sums over atomic contributions. PAINN is trained on 110k examples while 10k molecules are used as a validation set for decaying the learning rate and early stopping. The remaining data is used as test set and results are averaged over three random splits. For the isotropic polarizability  $\alpha$ , we first observed validation MAEs of  $0.054 a_0$ . Upon closer inspection, we notice that for this property both the squared loss as well as the MAE can be reduced when minimizing the MAE directly (as done by Klicpera et al. (2020b)). This yields both validation and test MAEs of  $0.045 a_0$  that are comparable to those of DIMENET++.

Tab. 2 shows the mean absolute error (MAE) of PAINN for 12 target properties of QM9 in comparison with previous approaches. SCHNET (Schütt et al., 2017) and PHYSNET (Unke & Meuwly, 2019) are MPNNs with distance-based interactions, DIMENET++ (Klicpera et al., 2020b) includes additional angular information and is an improved variant of DIMENET (Klicpera et al., 2020a). LINET (Miller et al., 2020) and CORMORANT (Anderson et al., 2019) are equivariant neural networks based on spherical harmonics and Clebsch-Gordon coefficients.

PAINN achieves state-of-the-art results in six target properties and yields comparable results to DIMENET++ on another two targets. On the remaining properties, PAINN achieves the second best results after DIMENET++. Note that PAINN using about 600k parameters is significantly smaller than DIMENET++ with about 1.8M parameters. For random batches of 50 molecules from QM9, the inference time is reduced from 45 ms to 13 ms, i.e. an improvement of more than 70%, when comparing PAINN to the reference implementation of DIMENET++<sup>1</sup> using an NVIDIA V100.

### 5.2. Molecular dynamics trajectories

We evaluate the ability to predict combined energies and forces on the MD17 benchmark (Chmiela et al., 2017) including molecular dynamics trajectories of small organic molecules. While the atomic forces could be predicted directly from vectorial features, we employ the gradients of the energy model  $\vec{F}_i = -\partial E / \partial \vec{r}_i$  to ensure conservation of energy. This property is crucial to run stable molecular dynamics simulations. To demonstrate the data efficiency of PAINN, we use the more challenging setting with 1k known structures of which we use 950 for training and 50 for validation, where a separate model is trained for each trajectory. Tab. 3 shows the comparison with sGDML (Chmiela et al., 2018) and NEQUIP (Batzner et al., 2021), which were trained on forces only, as well as SCHNET, PHYSNET, DIMENET and FCHL19 (Christensen et al., 2020), that were trained on a combined loss of energies and forces. Christensen & von Lilienfeld (2020) have found that the en-

<sup>1</sup><https://github.com/klicperaj/dimenet>

Table 2. Mean absolute errors on QM9 dataset for various chemical properties. Results for PAiNN are averaged over three random splits. Best in **bold**.

Target	Unit	SCHNET	PHYSNET	DIMENET++	CORMORANT	LINET	PAiNN
	D	0.033	0.053	0.030	0.038	0.043	<b>0.012</b>
	$a_0^3$	0.235	0.062	<b>0.044</b>	0.085	0.088	0.045
HOMO	meV	41	32.9	<b>24.6</b>	34	46.0	27.6
LUMO	meV	34	24.7	<b>19.5</b>	38	34.6	20.4
	meV	63	42.5	<b>32.6</b>	38	67.5	45.7
$\hbar R^2 i$	$a_0^2$	0.073	0.765	0.331	0.961	0.354	<b>0.066</b>
ZPVE	meV	1.7	1.39	<b>1.21</b>	2.03	1.56	1.28
U <sub>O</sub>	meV	14	8.15	6.32	22	13.46	<b>5.85</b>
U	meV	19	8.34	6.28	21	13.83	<b>5.83</b>
H	meV	14	8.42	6.53	21	14.36	<b>5.98</b>
G	meV	14	9.40	7.56	20	13.99	<b>7.35</b>
$c_v$	$\frac{\text{cal}}{\text{mol K}}$	0.033	0.028	<b>0.023</b>	0.026	0.031	0.024

Table 3. Mean absolute errors on MD17 dataset for energy and force predictions in kcal/mol and kcal/mol/Å, respectively. Batzner et al. (2021) only reported force errors for NEQUIP. Results for PAiNN are averaged over three random splits. Best in **bold**.

		sGDML	NEQUIP	PAiNN	SCHNET	PHYSNET	DIMENET	FCHL19	PAiNN
		trained on forces only			trained on energies & forces				
ASPIRIN	energy	0.19	–	<b>0.167</b>	0.37	0.230	0.204	0.182	<b>0.159</b>
	forces	0.68	0.348	<b>0.338</b>	1.35	0.605	0.499	0.478	<b>0.371</b>
ETHANOL	energy	0.07	–	<b>0.064</b>	0.08	0.059	0.064	<b>0.054</b>	0.063
	forces	0.33	<b>0.208</b>	0.224	0.39	0.160	0.230	<b>0.136</b>	0.230
MALONDIALDEHYDE	energy	<b>0.10</b>	–	<b>0.100</b>	0.13	0.094	0.104	<b>0.081</b>	0.091
	forces	0.41	<b>0.337</b>	0.344	0.66	0.319	0.383	<b>0.245</b>	0.319
NAPHTHALENE	energy	<b>0.12</b>	–	<b>0.116</b>	0.16	0.142	0.122	<b>0.117</b>	<b>0.117</b>
	forces	0.11	0.096	<b>0.077</b>	0.58	0.310	0.215	0.151	<b>0.083</b>
SALICYLIC ACID	energy	<b>0.12</b>	–	<b>0.116</b>	0.20	0.126	0.134	<b>0.114</b>	<b>0.114</b>
	forces	0.28	0.238	<b>0.195</b>	0.85	0.337	0.374	0.221	<b>0.209</b>
TOLUENE	energy	<b>0.10</b>	–	<b>0.095</b>	0.12	0.100	0.102	0.098	<b>0.097</b>
	forces	0.14	0.101	<b>0.094</b>	0.57	0.191	0.216	0.203	<b>0.102</b>
URACIL	energy	<b>0.11</b>	–	<b>0.106</b>	0.14	0.108	0.115	<b>0.104</b>	<b>0.104</b>
	forces	0.24	0.172	<b>0.139</b>	0.56	0.218	0.301	<b>0.105</b>	0.140

ergies of MD17 are noisy. Thus, depending on the molecule and chosen tradeoff, using energies for training is not always beneficial. For this reason, we train two PAiNN models per trajectory: only on forces and on a combined loss including energy with a force error weight of  $\rho = 0.95$ . PAiNN achieves the lowest mean absolute errors for 12 out of 14 targets on models trained only on forces and exhibits errors in a similar range as Gaussian regression with the FCHL19 kernel. Overall, PAiNN performs best or equal to FCHL19 on 9 out of 14 targets. This demonstrates that equivariant neural networks approaches such as PAiNN are able to compete with kernel methods in the small data regime, while being able to scale to large data sets at the same time.

### 5.3. Advantages of equivariant features

#### 5.3.1. ABLATION STUDIES

We evaluate the impact of equivariant vector features at the example of the aspirin MD trajectory from the previous section. Compared to the full model, we remove the scalar product of vector features in Eq. 9 from the update block and the convolution over vector features in Eq. 8 (i.e.,  $W_{VV} = 0$ ).

Table 4. Ablation study for the prediction of energies [kcal/mol] and forces [kcal/mol/Å] for aspirin trajectories from MD17.

Ablation	# params	$F$	energy MAE	force MAE
no ablation	588.3k	128	0.159	0.371
no scalar product of vector features in Eq. 8	589.1k	134	0.173	0.420
no vector propagation ( $W_{VV} = 0$ in Eq. 7)	589.2k	135	0.183	0.441
remove both	590.1k	142	0.200	0.507
no vector features	590.3k	174	0.449	1.194

Table 4 show the results for the various ablations. The number of parameters is kept approximately constant by raising the number of node features  $F$  accordingly. We observe that all ablated components contribute to the final accuracy of the model, where the convolution over equivariant features in the message function has a slightly larger impact. This component also enables the propagation of directional information, which will be examined in Sec. 5.3.2. Finally, we remove all vector features from the model, resulting in an invariant model. Despite keeping the number of param-




Figure 4. Rotational energy profile of substituted ferrocene obtained by varying the rotation angle  $\theta$  while keeping all bond distances fixed at their equilibrium values (H: white, C: black, F: purple, Fe: grey). For small cutoffs  $r_{\text{cut}}$  (dashed grey circles), MPNNs with scalar feature representations are unable to represent information about  $\theta$ .

eters constant by increasing the number of atoms features to  $F = 174$ , the mean absolute error of the forces increases beyond 1 kcal/mol/Å.

### 5.3.2. PROPAGATION OF DIRECTIONAL INFORMATION IN SUBSTITUTED FERROCENE

To demonstrate the advantages of equivariant over invariant representations in practice, we consider a ferrocene derivative where one hydrogen atom in each cyclopentadienyl ring is substituted by fluorine (see Fig. 4). This molecule has been chosen as it features small energy fluctuations ( $<1$  kcal/mol) when the rings rotate relative to each other. Since the torsional energy profile depends mainly on the orientation of the distant fluorine atoms (measured by the rotation angle  $\theta$ ), it is a challenging prediction target for models without equivariant representations. Fig. 4 shows the predicted energy profiles for a full rotation of the cyclopentadienyl ring using cutoffs  $r_{\text{cut}} \in \{2.5, 3.0, 4.0\}$  Å. All models are trained on energies of 10k structures sampling thermal fluctuations (300 K) and ring rotations of substituted ferrocene. Another 1k structures are used for validation. SchNet (Schütt et al., 2017), which uses scalar features, predicts a flat energy profile for  $r_{\text{cut}} = 3.0$  Å, because it is unable to resolve  $\theta$ . Although it is possible to encode the relevant information when using a larger cutoff, the learned representation seems to be susceptible to noise, as indicated by the deviations of the predicted energy profile. Since DimeNet (Klicpera et al., 2020a) includes terms that explicitly depend on angles between triplets of atoms,  $\theta$  can be resolved for smaller cutoffs, but still fails

for  $r_{\text{cut}} = 2.5$  Å. In contrast, the equivariant environment representation of PAiNN allows to faithfully reproduce the torsional energy profile even for very small cutoffs.

## 5.4. Molecular Spectra

We apply PAiNN to the efficient computation of infrared and Raman spectra of ethanol and aspirin. Although these spectra can in principle be obtained from a single molecular structure via the harmonic oscillator approximation, such an approach tends to neglect important effects, e.g. different molecular conformations (Thomas et al., 2013). In order to obtain high quality spectra, molecular dynamics simulations have to be performed, requiring the prediction of forces at each time step. Additionally, one needs to compute dipole moments and polarizabilities along the generated trajectory. The infrared and Raman spectra are obtained as the Fourier transform of the time auto-correlation functions of the respective property. Using electronic structure methods, such simulations become prohibitively expensive due to the large number of successive computations required. Since nuclear quantum effects (NQE) need to be considered to obtain high quality spectra (Sauceda et al., 2021), ring-polymer molecular dynamics (RPMD) simulations need to be performed, which treat multiple copies of a molecule simultaneously and further increase computational cost.

We train a joined model for energies, forces, dipole moments and polarizability tensors on 8k ethanol conformations using additional 1k molecules each as validation and test set. A second model is trained for aspirin with 16k training examples and 2k molecules for validation and testing. Energies and forces are predicted as described above for MD17. The dipole moments and polarizability tensors are obtained as described in Section 4.2, employing two gated equivariant blocks each yielding atomwise scalars and vectors, that are used to compute the outputs as described in Eqs. 13 and 14. The joined model exhibits accurate predictions for energy, forces, dipole moments and polarizabilities, as shown in the inset table of Fig 5. To evaluate the importance of equivariant features for the dipole prediction, we have trained an alternative model using only latent partial charges, but no local dipoles. The results show that predicting molecular dipoles using vector features (Eq. 13) reduces the mean absolute error by more than 50%, reaching 0.009 D compared to 0.020 D using only scalars and positions (Eq. 12). Simulated infrared and Raman spectra using classical MD and RPMD with 64 replicas and the respective PAiNN models are shown for both molecules in Fig. 5, alongside harmonic spectra obtained with the electronic structure reference and, for ethanol, experimental spectra recorded in the gas phase (Linstrom & Mallard, 2020; Kiefer, 2017).

The peak positions and intensities of ethanol infrared spectra (Fig. 5a, top) computed with classical MD and PAiNN



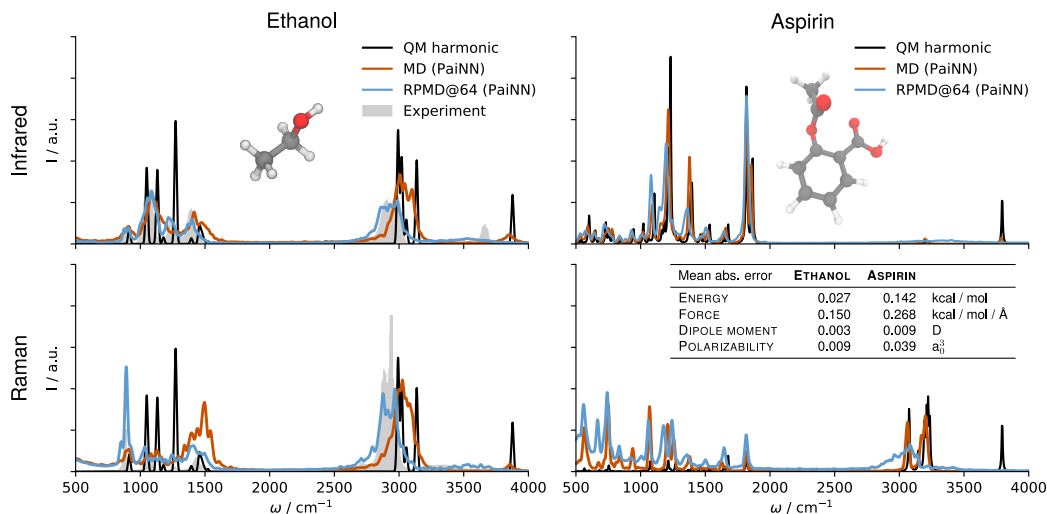


Figure 5. IR (top) and Raman (bottom) spectra of ethanol and aspirin. Spectra calculated with the reference method using the harmonic oscillator approximation are shown in black (QM harmonic). The inset table shows the mean absolute errors on the respective test set.

agree closely with the static electronic structure reference, e.g. in the C-H and O-H stretching regions at  $3000\text{ cm}^{-1}$  and  $3900\text{ cm}^{-1}$ . This indicates, that PAINN faithfully reproduces the original electronic structure method in general. However, when compared to experiment, both spectra are shifted towards higher frequencies. The PAINN RMPD spectrum, on the other hand shows excellent agreement with experiment, illustrating the importance of NQEs. Similar observations can be made for the ethanol Raman spectrum (Fig. 5a, bottom), where the PAINN RMPD spectrum once again offers the most faithful reproduction of the experiment. We observe similar trends for aspirin (Fig. 5b). RMPD spectra are once again shifted to lower frequencies while the classical MD closely resembles the static electronic structure spectrum.

PAINN reduces the time required to obtain these spectra by several orders of magnitude. For ethanol, an individual reference computation takes 140 seconds compared to 14 ms for a PAINN evaluation on a V100 GPU. This corresponds to a reduction of the overall simulation time from 400 days to approximately one hour. The speedup offered by PAINN is even more pronounced for aspirin, where a simulation that would have taken 25 years (3140 seconds / step) now takes one hour (15 ms / step).

## 6. Conclusions

We have given general guidelines to design equivariant MPNNs and discussed the advantages of equivariant representations over angular features in terms of computational efficiency as well as their ability to propagate directional information. On this basis, we have proposed PAINN that yields fast and accurate predictions of scalar and tensorial

molecular properties. Thereby, equivariant message passing allows us to significantly reduce both model size and inference time compared to directional message-passing while retaining accuracy. Finally, we have demonstrated that PAINN can be applied to the prediction of tensorial properties, which we leverage to accelerate the simulation of molecular spectra by 4-5 orders of magnitude – from years to hours.

In future work, the equivariant representation of PAINN as well as the ability to predict tensorial properties may be leveraged in generative models of 3d geometries (Gebauer et al., 2019; Köhler et al., 2019; Simm et al., 2020) or the prediction of wavefunctions (Hegde & Bowen, 2017; Schütt et al., 2019; Hermann et al., 2020). We see further applications of equivariant message passing in 3d shape recognition and graph embedding (Goyal & Ferrara, 2018).

Many challenges remain for the fast and accurate prediction of molecular properties, e.g. modeling of enzymatic active sites or surface reactions. To describe such phenomena, highly accurate reference methods are required. Due to their computational cost, reference data generation can become a bottleneck, making data-efficient MPNNs such as PAINN invaluable for future chemistry research.

## Acknowledgements

KTS and OTU acknowledge support by the Federal Ministry of Education and Research (BMBF) for BIFOLD (01IS18037A). OTU acknowledges funding from the Swiss National Science Foundation (Grant No. P2BSP2\_188147). MG works at the BASLEARN – TU Berlin/BASF Joint Lab for Machine Learning, co-financed by TU Berlin and BASF SE.

## References

- Anderson, B., Hy, T. S., and Kondor, R. Cormorant: Covariant Molecular Neural Networks. *Advances in Neural Information Processing Systems*, 32:14537–14546, 2019.
- Bartók, A. P., Payne, M. C., Kondor, R., and Csányi, G. Gaussian approximation potentials: The accuracy of quantum mechanics, without the electrons. *Phys. Rev. Lett.*, 104(13):136403, 2010.
- Bartók, A. P., Kondor, R., and Csányi, G. On representing chemical environments. *Phys. Rev. B*, 87(18):184115, 2013.
- Bartók, A. P., Kermode, J., Bernstein, N., and Csányi, G. Machine learning a general-purpose interatomic potential for silicon. *Phys. Rev. X*, 8(4):041048, 2018.
- Batzner, S., Smidt, T. E., Sun, L., Mailoa, J. P., Kornbluth, M., Molinari, N., and Kozinsky, B. SE (3)-Equivariant Graph Neural Networks for Data-Efficient and Accurate Interatomic Potentials. *arXiv preprint arXiv:2101.03164*, 2021.
- Behler, J. Atom-centered symmetry functions for constructing high-dimensional neural network potentials. *J. Chem. Phys.*, 134(7):074106, 2011.
- Behler, J. Perspective: Machine learning potentials for atomistic simulations. *J. Chem. Phys.*, 145(17):170901, 2016.
- Behler, J. and Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.*, 98(14):146401, 2007.
- Chmiela, S., Tkatchenko, A., Sauceda, H. E., Poltavsky, I., Schütt, K. T., and Müller, K.-R. Machine Learning of Accurate Energy-Conserving Molecular Force Fields. *Sci. Adv.*, 3(5):e1603015, 2017.
- Chmiela, S., Sauceda, H. E., Müller, K.-R., and Tkatchenko, A. Towards exact molecular dynamics simulations with machine-learned force fields. *Nat. Commun.*, 9(1):3887, 2018. doi: 10.1038/s41467-018-06169-2.
- Christensen, A. S. and von Lilienfeld, O. A. On the role of gradients for machine learning of molecular energies and forces. *Mach. Learn.: Sci. Technol.*, 1(4):045018, 2020.
- Christensen, A. S., Bratholm, L. A., Faber, F. A., and Anatole von Lilienfeld, O. FCHL revisited: Faster and more accurate quantum machine learning. *J. Chem. Phys.*, 152(4):044107, 2020.
- Cohen, T. S. and Welling, M. Steerable CNNs. *International Conference on Learning Representations*, 2017.
- Duvenaud, D. K., Maclaurin, D., Iparraguirre, J., Bombarell, R., Hirzel, T., Aspuru-Guzik, A., and Adams, R. P. Convolutional networks on graphs for learning molecular fingerprints. *Advances in Neural Information Processing Systems*, 28:2224–2232, 2015.
- Gastegger, M., Behler, J., and Marquetand, P. Machine learning molecular dynamics for the simulation of infrared spectra. *Chem. Sci.*, 8(10):6924–6935, 2017.
- Gastegger, M., Schütt, K. T., and Müller, K.-R. Machine learning of solvent effects on molecular spectra and reactions. *arXiv preprint arXiv:2010.14942*, 2020.
- Gebauer, N., Gastegger, M., and Schütt, K. Symmetry-adapted generation of 3d point sets for the targeted discovery of molecules. *Advances in Neural Information Processing Systems*, pp. 7564–7576, 2019.
- Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. Neural message passing for quantum chemistry. *International Conference on Machine Learning*, 2017.
- Goyal, P. and Ferrara, E. Graph embedding techniques, applications, and performance: A survey. *Knowledge-Based Systems*, 151:78–94, 2018.
- Hegde, G. and Bowen, R. C. Machine-learned approximations to Density Functional Theory Hamiltonians. *Sci. Rep.*, 7(1):1–11, 2017.
- Hermann, A., Krawczyk, R. P., Lein, M., Schwerdtfeger, P., Hamilton, I. P., and Stewart, J. J. Convergence of the many-body expansion of interaction potentials: From van der waals to covalent and metallic systems. *Physical Review A*, 76(1):013202, 2007.
- Hermann, J., Schätzle, Z., and Noé, F. Deep-neural-network solution of the electronic Schrödinger equation. *Nat. Chem.*, 12(10):891–897, 2020.
- Hinton, G. E., Krizhevsky, A., and Wang, S. D. Transforming auto-encoders. *International Conference on Artificial Neural Networks*, pp. 44–51, 2011.
- Jing, B., Eismann, S., Suriana, P., Townshend, R. J. L., and Dror, R. Learning from protein structure with geometric vector perceptrons. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=1YLJDvSx6J4>.
- Kearnes, S., McCloskey, K., Berndl, M., Pande, V., and Riley, P. Molecular graph convolutions: moving beyond fingerprints. *J. Comput. Aided Mol. Des.*, 30(8):595–608, 2016.

- Kiefer, J. Simultaneous acquisition of the polarized and depolarized Raman signal with a single detector. *Anal. Chem.*, 89(11):5725–5728, 2017.
- Kingma, D. P. and Ba, J. Adam: A method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Klicpera, J., Groß, J., and Günnemann, S. Directional message passing for molecular graphs. *International Conference on Learning Representations*, 2020a.
- Klicpera, J., Groß, J., and Günnemann, S. Fast and Uncertainty-Aware Directional Message Passing for Non-Equilibrium Molecules. *Machine Learning for Molecules Workshop at NeurIPS*, 2020b.
- Köhler, J., Klein, L., and Noé, F. Equivariant Flows: sampling configurations for multi-body systems with symmetric energies. *Proceedings of the 37th International Conference on Machine Learning*, 2019.
- Kondor, R. and Trivedi, S. On the Generalization of Equivariance and Convolution in Neural Networks to the Action of Compact Groups. *International Conference on Machine Learning*, pp. 2747–2755, 2018.
- Linstrom, P. and Mallard, W. G. *NIST Chemistry WebBook NIST Standard Reference Database Number 69*. National Institute of Standards and Technology, doi:10.18434/T4D303, (retrieved September 24, 2020), Gaithersburg MD, 20899, 2020.
- Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. *International Conference on Learning Representations*, 2019.
- Lu, D., Wang, H., Chen, M., Liu, J., Lin, L., Car, R., Jia, W., Zhang, L., et al. 86 PFLOPS Deep Potential Molecular Dynamics simulation of 100 million atoms with ab initio accuracy. *arXiv preprint arXiv:2004.11658*, 2020.
- Lubbers, N., Smith, J. S., and Barros, K. Hierarchical modeling of molecular energies using a deep neural network. *J. Chem. Phys.*, 148(24):241715, 2018.
- Miller, B. K., Geiger, M., Smidt, T. E., and Noé, F. Relevance of rotationally equivariant convolutions for predicting molecular properties. *arXiv preprint arXiv:2008.08461*, 2020.
- Morawietz, T., Singraber, A., Dellago, C., and Behler, J. How van der Waals interactions determine the unique properties of water. *Proc. Natl. Acad. Sci.*, 113(30):8368–8373, 2016.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint arXiv:1912.01703*, 2019.
- Pozdnyakov, S. N., Willatt, M. J., Bartók, A. P., Ortner, C., Csányi, G., and Ceriotti, M. Incompleteness of atomic structure representations. *Physical Review Letters*, 125(16):166001, 2020.
- Ramakrishnan, R., Dral, P. O., Rupp, M., and Von Lilienfeld, O. A. Quantum chemistry structures and properties of 134 kilo molecules. *Sci. Data*, 1(1):1–7, 2014.
- Sauceda, H. E., Vassilev-Galindo, V., Chmiela, S., Müller, K.-R., and Tkatchenko, A. Dynamical strengthening of covalent and non-covalent molecular interactions by nuclear quantum effects at finite temperature. *Nat. Commun.*, 12(1):1–10, 2021.
- Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., and Monfardini, G. The graph neural network model. *IEEE Trans. Neural Netw.*, 20(1):61–80, 2008.
- Schütt, K., Kindermans, P.-J., Sauceda, H. E., Chmiela, S., Tkatchenko, A., and Müller, K.-R. SchNet: A continuous-filter convolutional neural network for modeling quantum interactions. *Advances in Neural Information Processing Systems*, pp. 991–1001, 2017.
- Schütt, K., Kessel, P., Gastegger, M., Nicoli, K., Tkatchenko, A., and Müller, K.-R. SchNetPack: A deep learning toolbox for atomistic systems. *J. Chem. Theory Comput.*, 15(1):448–455, 2018.
- Schütt, K., Gastegger, M., Tkatchenko, A., Müller, K.-R., and Maurer, R. J. Unifying machine learning and quantum chemistry with a deep neural network for molecular wavefunctions. *Nat. Commun.*, 10(1):1–10, 2019.
- Schütt, K. T., Arbabzadah, F., Chmiela, S., Müller, K. R., and Tkatchenko, A. Quantum-chemical insights from deep tensor neural networks. *Nat. Commun.*, 8, 2017. doi: 10.1038/ncomms13890.
- Schütt, K. T., Sauceda, H. E., Kindermans, P.-J., Tkatchenko, A., and Müller, K.-R. SchNet—A deep learning architecture for molecules and materials. *J. Chem. Phys.*, 148(24):241722, 2018.
- Simm, G. N., Pinsler, R., Csányi, G., and Hernández-Lobato, J. M. Symmetry-Aware Actor-Critic for 3D Molecular Design. *arXiv preprint arXiv:2011.12747*, 2020.
- Thomas, M., Brehm, M., Fligg, R., Vöhringer, P., and Kirchner, B. Computing vibrational spectra from ab initio molecular dynamics. *Phys. Chem. Chem. Phys.*, 15(18): 6608–6622, 2013.

- Thomas, N., Smidt, T., Kearnes, S., Yang, L., Li, L., Kohlhoff, K., and Riley, P. Tensor field networks: Rotation-and translation-equivariant neural networks for 3D point clouds. *arXiv preprint arXiv:1802.08219*, 2018.
- Unke, O. T. and Muwly, M. PhysNet: A Neural Network for Predicting Energies, Forces, Dipole Moments, and Partial Charges. *J. Chem. Theory Comput.*, 15(6):3678–3693, 2019.
- Unke, O. T., Chmiela, S., Sauceda, H. E., Gastegger, M., Poltavsky, I., Schütt, K. T., Tkatchenko, A., and Müller, K.-R. Machine Learning Force Fields. *arXiv preprint arXiv:2010.07067*, 2020.
- Veit, M., Wilkins, D. M., Yang, Y., DiStasio Jr, R. A., and Ceriotti, M. Predicting molecular dipole moments by combining atomic partial charges and atomic dipoles. *J. Chem. Phys.*, 153(2):024113, 2020.
- von Lilienfeld, O. A., Müller, K.-R., and Tkatchenko, A. Exploring chemical compound space with quantum-based machine learning. *Nat. Rev. Chem.*, 4(7):347–358, 2020.
- Weiler, M., Geiger, M., Welling, M., Boomsma, W., and Cohen, T. S. 3D Steerable CNNs: Learning Rotationally Equivariant Features in Volumetric Data. *Advances in Neural Information Processing Systems*, 31:10381–10392, 2018a.
- Weiler, M., Hamprecht, F. A., and Storath, M. Learning Steerable Filters for Rotation Equivariant CNNs. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 849–858, 2018b.
- Westermayr, J., Gastegger, M., and Marquetand, P. Combining SchNet and SHARC: The SchNarc machine learning approach for excited-state dynamics. *J. Phys. Chem. Lett.*, 11(10):3828–3834, 2020.
- Worrall, D. and Brostow, G. CubeNet: Equivariance to 3D Rotation and Translation. *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 567–584, 2018.