

Magnitude-Modulated Equivariant Adapter for Parameter-Efficient Fine-Tuning of Equivariant Graph Neural Networks

Dian Jin¹, Yancheng Yuan^{2*}, Xiaoming Tao¹

¹Research Institute for Intelligent Wearable Systems, The Hong Kong Polytechnic University

²Department of Applied Mathematics, The Hong Kong Polytechnic University

diann.jin@connect.polyu.hk, yancheng.yuan@polyu.edu.hk

Abstract

Pretrained equivariant graph neural networks based on spherical harmonics offer efficient and accurate alternatives to computationally expensive ab-initio methods, yet adapting them to new tasks and chemical environments still requires fine-tuning. Conventional parameter-efficient fine-tuning (PEFT) techniques, such as Adapters and LoRA, typically break symmetry, making them incompatible with those equivariant architectures. ELoRA, recently proposed, is the first equivariant PEFT method. It achieves improved parameter efficiency and performance on many benchmarks. However, the relatively high degrees of freedom it retains within each tensor order can still perturb pretrained feature distributions and ultimately degrade performance. To address this, we present Magnitude-Modulated Equivariant Adapter (MMEA), a novel equivariant fine-tuning method which employs lightweight scalar gating to modulate feature magnitudes on a per-order and per-multiplicity basis. We demonstrate that MMEA preserves strict equivariance and, across multiple benchmarks, consistently improves energy and force predictions to state-of-the-art levels while training fewer parameters than competing approaches. These results suggest that, in many practical scenarios, modulating channel magnitudes is sufficient to adapt equivariant models to new chemical environments without breaking symmetry, pointing toward a new paradigm for equivariant PEFT design.

Code — <https://github.com/CLaSLoVe/MMEA>

Extended version — <https://arxiv.org/pdf/2511.06696>

1 Introduction

Density-functional theory (DFT) offers a reliable description of electronic structure, and has already become a standard technique in most branches of chemistry and materials science (Burke 2012). However, even with an exascale system, the scale one could achieve with DFT is limited due to its cubic scaling on system size (Fiedler et al. 2023). Thus, the costs of considering a huge number of configurations in long simulations, or of dealing with the exponentially growing number of possible materials in high-throughput screenings will remain prohibitive (Behler 2017).

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

To bridge this gap, researchers have turned to deep learning-based potentials (Behler and Parrinello 2007), which greatly accelerate computer simulations, while preserving quantum mechanical accuracy (Behler 2017). Those methods (Behler and Parrinello 2007; Smith, Isayev, and Roitberg 2017; Schütt et al. 2018; Gilmer et al. 2017; Zhang et al. 2018; Unke and Meuwly 2019; Unke et al. 2021; Aykent and Xia 2025; Chang and Zhu 2025), hold significant promise in addressing the accuracy-efficiency trade-off and therefore allow the expansion of the spatial and temporal scales of molecular dynamics simulations (Wang et al. 2025). Recent work has produced large, pretrained foundation models such as Equiformer (Liao and Smidt 2022), NequIP (Batzner et al. 2022), Uni-Mol (Zhou et al. 2023), MACE (Batatia et al. 2022). Among all these models, those equivariant graph neural networks based on spherical harmonics are particularly powerful, as they are not only inherently designed to respect rotational, translational, and permutational symmetries, but also to model high-order physical information (Anderson, Hy, and Kondor 2019). This leads to exceptional sample efficiency: chemical accuracy can often be achieved with just a few hundred to a few thousand local structures, as it is an important benefit when high-quality quantum data is costly to obtain (Maruf, Kim, and Ahmad 2025).

Nonetheless, when a target system is poorly represented in the pretraining data, such as rare chemistry configurations, a foundation model can lose its accuracy. Fortunately, it has been shown in multiple domains (Howard and Ruder 2018; Devi, Butler, and Sai Gautam 2024) that fine-tuning on a small, task-specific dataset can restore accuracy in such cases and typically yields better results than training from scratch. However, full-parameter fine-tuning risks over-fitting and catastrophic forgetting (Bethune et al. 2025). Parameter-efficient fine-tuning (PEFT) techniques like Low-Rank Adaptation (LoRA) (Hu et al. 2022) are attractive alternatives, yet mixing different tensor orders in an equivariant network breaks symmetry.

An equivariant extension of LoRA called ELoRA (Wang et al. 2025), addresses this issue by introducing path-dependent low-rank adapters into each tensor channel, demonstrating superior performance compared to full-parameter fine-tuning. This clearly highlights the potential of PEFT. By restricting the degrees of freedom during learn-

ing, ELoRA reduces the risk of overfitting and catastrophic forgetting, resulting in a model with stronger generalization capabilities.

However, we observe that ELoRA still maintains relatively high degrees of freedom within each tensor order. Building on the insight that a well-trained equivariant backbone inherently provides a robust foundation for each tensor order, we propose an even lighter fine-tuning strategy that employs only per-channel scaling adjustments. This approach strictly preserves equivariance, introduces fewer additional parameters, and consistently outperforms ELoRA across diverse molecular and materials benchmarks.

Our main contributions are listed below:

1. We introduce MMEA, an equivariant PEFT approach for equivariant graph neural networks based on spherical harmonics that leverages dynamic magnitude modulation on a per-channel and per-multiplicity basis.
2. Evaluations on multiple molecular benchmarks show that MMEA achieves state-of-the-art fine-tuning performance, surpassing previous methods while requiring fewer trainable parameters.
3. We integrate MMEA into the widely used e3nn framework, providing a convenient and efficient fine-tuning solution for the community.

2 Related Works

2.1 Equivariant Graph Neural Networks

Graph Neural Networks (GNNs), especially those based on the Message Passing Neural Network (MPNN) framework (Gilmer et al. 2017), have demonstrated strong performance in processing graph-structured data such as social networks (Liu et al. 2021) and recommender systems (Wu et al. 2022). In molecular applications, small molecules can naturally be represented as graphs, where atoms serve as nodes and bonds as edges. Consequently, GNNs have been widely adopted for molecular property prediction (Lee et al. 2023; Wieder et al. 2020). However, unlike abstract graphs, molecules possess three-dimensional structures that are crucial for determining their physical and chemical properties, such as energy and atomic forces. This motivates the integration of 3D structural information into GNN models.

Early efforts focused on incorporating invariant geometric features into GNNs to respect the rotational and translational symmetries of molecular systems. For example, SchNet (Schütt et al. 2018) encodes pairwise atomic distances (i.e., two-body information) using radial basis functions. DimeNet (Gasteiger, Groß, and Günnemann 2020) extends this approach by modeling angular information between triplets of atoms (three-body interactions), while GemNet (Gasteiger, Becker, and Günnemann 2021) further includes torsional angles (four-body interactions). These models, often referred to as invariant GNNs, produce outputs that remain unchanged under global rotations or translations of the molecule. While effective, they rely on carefully designed geometric descriptors and handcrafted feature engineering, which limit scalability and make it difficult to model directional quantities such as forces.

To overcome these limitations, researchers have proposed equivariant GNNs, whose outputs transform predictably under geometric transformations. These models can capture both scalar and directional information and are generally more expressive. Two major classes of equivariant GNNs have emerged.

The first class includes scalar–vector models such as PaiNN (Schütt, Unke, and Gastegger 2021), EGNN (Satorras, Hoogeboom, and Welling 2021), and TorchMD-Net (Thölke and De Fabritiis 2022). These methods maintain both scalar and three-dimensional vector features at each node, updating them jointly during message passing. They are computationally efficient and suitable for tasks involving directionally dependent quantities. However, from a theoretical perspective, these models only capture first-order tensor interactions and thus exhibiting limited fitting capability. Moreover, some studies have shown that such architectures may suffer from the “zero function” problem (Cen et al. 2024), where outputs can collapse to zero under certain conditions, limiting their expressiveness.

The second class of methods involves higher-order tensor models based on spherical harmonics and Clebsch–Gordan (CG) tensor products. Tensor Field Net (Thomas et al. 2018) and Cormorant (Anderson, Hy, and Kondor 2019) pioneered this strategy by decomposing features into spherical harmonic components and using CG tensor products to achieve equivariant message passing. These higher-order tensor-based equivariant models excel primarily because they can explicitly encode higher-order geometric and physical information, thereby capturing complex interactions that standard GNNs find difficult to represent. In particular, the subsequent MACE family of foundation models (Bataatia et al. 2022) extends the architecture from traditional pairwise interactions to higher-order many-body interactions, significantly improving the accuracy of energy and force predictions.

2.2 Parameter-Efficient Fine-Tuning

PEFT methods aim to reduce the number of parameters that need to be updated during the fine-tuning process. The main goal is to address the inefficiency of full-parameter updates as model size and task complexity increase. By limiting the degrees of freedom in model updates, PEFT methods can help mitigate overfitting in certain scenarios. PEFT is primarily applied in CV and NLP, and there are many different approaches. For example, Adapter (Houlsby et al. 2019) tuning introduces bottleneck modules between layers to enable efficient adaptation; LoRA (Hu et al. 2022) decomposes weight matrices into low-rank representations, reducing the number of trainable parameters; BitFit (Zaken, Ravfogel, and Goldberg 2021) only updates bias terms while keeping all other parameters fixed; FiLM (Perez et al. 2018) modifies intermediate representations through feature-wise affine transformations, enabling task-specific modulation without full parameter updates. Recently, research has extended these techniques to other domains, such as GNNs, including methods like AdapterGNN (Li, Han, and Bai 2024). However, most existing approaches map inputs to bottleneck modules and then back to outputs, which mixes ir-

reducible representations of different orders in equivariant GNNs. This breaks the equivariance property and leads to the loss of symmetry guarantees, which can negatively impact performance (Wang et al. 2025). To address this, ELoRA (Wang et al. 2025) is proposed, which introduces a path-dependent decomposition for weight updates in $SO(3)$ equivariant GNNs, ensuring that low-rank adaptations respect the symmetry constraints intrinsic to the model. This not only reduces the number of trainable parameters, but also enhanced prediction performance, enhancing data efficiency.

3 Methodology

We propose a novel equivariant adapter, Magnitude-Modulated Equivariant Adapter, for the fine-tuning process of Equivariant GNNs. Our approach is grounded in the path-dependent paradigm, but unlike the LoRA-based ELoRA method, MMEA employs scalar gating to adjust each channel’s magnitude without mixing different multiplicities under the same irreducible order. Consequently, it better leverages the information in the pretrained model, mitigates overfitting, and introduces fewer additional parameters.

The main architecture of MMEA is shown in Fig. 1. Our proposed method is inspired by physical intuition: in a well-trained equivariant GNN model, the multiplicity channel for each order already constitutes a robust basis, and allowing them to mix freely during fine-tuning may distort the geometry of the pretrained feature space. Our method avoids mixing different multiplicities, which helps preserve the knowledge learned by the pre-trained model more effectively. Next, we introduce the details of the MMEA architecture.

Node feature space The node feature space can be denoted as a direct sum of irreducible representations (Thomas et al. 2018):

$$\mathcal{H} := \bigoplus_{\ell=0}^L \mathcal{H}^{(\ell)}, \quad \mathcal{H}^{(\ell)} := V^{(\ell)} \otimes \mathbb{R}^{1 \times m_\ell}, \quad (1)$$

where $V^{(\ell)}$ is the irreducible representation of order ℓ with dimension $d_\ell = 2\ell + 1$, m_ℓ is its multiplicity, and \otimes denotes the Kronecker product. L is the maximum representation order. The $SO(3)$ (Special Orthogonal Group in 3 Dimensions) action g is

$$g \cdot (v \otimes a) := (\rho^{(\ell)}(g)v) \otimes a, \quad v \in V^{(\ell)}, \quad a \in \mathbb{R}^{1 \times m_\ell}, \quad (2)$$

with $\rho^{(\ell)} \in \mathbb{R}^{d_\ell \times d_\ell}$ is the Wigner- D matrix (Thomas et al. 2018) of order ℓ . One can realize that the group action g only acts on $V^{(\ell)}$, and the multiplicity space $\mathbb{R}^{1 \times m_\ell}$ is invariant.

A Lightweight gating network In MMEA, we only take $\mathbf{h}^{(0)}$ as input to maintain parameter efficiency. Then, we design a two-layer MLP to obtain multiplicity-wise gain scalars. In particular, it consists of a bottleneck projection

$$\mathbf{z} = \text{SiLU}(W_\downarrow \mathbf{h}^{(0)} + \mathbf{b}_\downarrow), \quad W_\downarrow \in \mathbb{R}^{r \times m_0}, \quad (3)$$

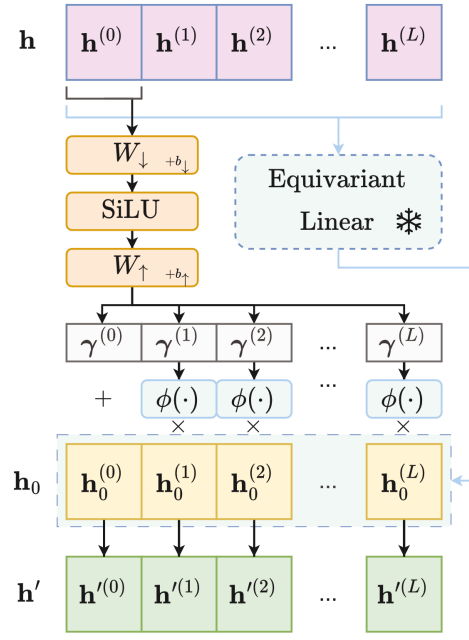


Figure 1: Overview of the Magnitude-Modulated Equivariant Adapter. Rounded boxes denote operators, dashed outline indicates components of the original architecture (untrainable components).

and followed by an expansion:

$$[\gamma^{(0)}, \gamma^{(1)}, \dots, \gamma^{(L)}] = W_\uparrow \mathbf{z} + \mathbf{b}_\uparrow \quad (4)$$

$$W_\uparrow \in \mathbb{R}^{(\sum_{\ell=0}^L m_\ell) \times r}, \quad (5)$$

where each $\gamma^{(\ell)} \in \mathbb{R}^{m_\ell}$ assigns one scalar gain per multiplicity copy of order ℓ . r is the bottleneck dimension used in the MLP. SiLU is Sigmoid Linear Unit (Elfving, Uchibe, and Doya 2018), and $W_\uparrow, W_\downarrow, \mathbf{b}_\uparrow, \mathbf{b}_\downarrow$ are trainable parameters.

Equivariant modulation Write $\mathbf{h}^{(\ell)} = \sum_{k=1}^{m_\ell} \mathbf{v}_k^{(\ell)} \otimes \mathbf{e}_k^{(\ell)}$ with $\mathbf{v}_k^{(\ell)} \in V^{(\ell)}$ and $\{\mathbf{e}_k^{(\ell)}\}$ the standard basis of \mathbb{R}^{m_ℓ} . Define

$$\mathcal{A}_\Gamma^{(0)}(\mathbf{h}^{(0)}) = \mathbf{h}^{(0)} + \gamma^{(0)}, \quad (6)$$

$$\mathcal{A}_\Gamma^{(\ell)}(\mathbf{h}^{(\ell)}) = \sum_{k=1}^{m_\ell} \phi(\gamma_k^{(\ell)}) \mathbf{v}_k^{(\ell)} \otimes \mathbf{e}_k^{(\ell)}, \quad \ell \geq 1, \quad (7)$$

where $\phi(x) = 1 + x$ (residual scaling) or $\phi(x) = e^x$ (positive scaling). Then we obtain

$$\mathcal{A}_\Gamma(\mathbf{h}) := \bigoplus_{\ell=0}^L \mathcal{A}_\Gamma^{(\ell)}(\mathbf{h}^{(\ell)}), \quad (8)$$

yielding the modulated feature $\mathbf{h}' := \mathcal{A}_\Gamma(\mathbf{h}) \in \mathcal{H}$.

3.1 The Equivariance Property of MMEA

Now, we will show the equivariance property of our designed MMEA adapter. In other words, we will show that, for every $g \in \text{SO}(3)$,

$$\mathcal{A}_\Gamma(g \cdot \mathbf{h}) = g \cdot \mathcal{A}_\Gamma(\mathbf{h}). \quad (9)$$

Step 1: The gains $\gamma^{(\ell)}$ are invariant. Recall that the gains $\gamma^{(\ell)}$ are produced solely from the scalar component $\mathbf{h}^{(0)} \in \mathbb{R}^{m_0}$ via the lightweight MLP:

$$\mathbf{z} = \text{SiLU}(W_\downarrow \mathbf{h}^{(0)} + \mathbf{b}_\downarrow), \quad (10)$$

$$[\gamma^{(0)}, \gamma^{(1)}, \dots, \gamma^{(L)}] = W_\uparrow \mathbf{z} + \mathbf{b}_\uparrow. \quad (11)$$

Since $\mathbf{h}^{(0)}$ transforms under the trivial irreducible representation ($\ell = 0$), we have

$$g \cdot \mathbf{h}^{(0)} = \mathbf{h}^{(0)} \quad \forall g \in \text{SO}(3). \quad (12)$$

Consequently, \mathbf{z} and hence every $\gamma^{(\ell)}$ are unchanged by the group action:

$$\gamma^{(\ell)}(g \cdot \mathbf{h}^{(0)}) = \gamma^{(\ell)}(\mathbf{h}^{(0)}), \quad \forall \ell. \quad (13)$$

Step 2: Each block $\mathcal{A}_\Gamma^{(\ell)}$ commutes with the group action.

Decompose the ℓ -th order feature as

$$\mathbf{h}^{(\ell)} = \sum_{k=1}^{m_\ell} \mathbf{v}_k^{(\ell)} \otimes \mathbf{e}_k^{(\ell)}, \quad \mathbf{v}_k^{(\ell)} \in V^{(\ell)}, \quad (14)$$

the group action g acts only on $V^{(\ell)}$:

$$g \cdot \mathbf{h}^{(\ell)} = \sum_{k=1}^{m_\ell} \rho^{(\ell)}(g) \mathbf{v}_k^{(\ell)} \otimes \mathbf{e}_k^{(\ell)}. \quad (15)$$

By definition, the adapter block is

$$\mathcal{A}_\Gamma^{(0)}(\mathbf{h}^{(0)}) = \mathbf{h}^{(0)} + \gamma^{(0)}, \quad (16)$$

$$\mathcal{A}_\Gamma^{(\ell)}(\mathbf{h}^{(\ell)}) = \sum_{k=1}^{m_\ell} \phi(\gamma_k^{(\ell)}) \mathbf{v}_k^{(\ell)} \otimes \mathbf{e}_k^{(\ell)}, \quad \ell \geq 1. \quad (17)$$

For $\ell \geq 1$, because the gains $\gamma_k^{(\ell)}$ are group-invariant, applying $\mathcal{A}_\Gamma^{(\ell)}$ to the transformed features yields

$$\mathcal{A}_\Gamma^{(\ell)}(g \cdot \mathbf{h}^{(\ell)}) = \sum_{k=1}^{m_\ell} \phi(\gamma_k^{(\ell)}) (\rho^{(\ell)}(g) \mathbf{v}_k^{(\ell)} \otimes \mathbf{e}_k^{(\ell)}) \quad (18)$$

$$= \rho^{(\ell)}(g) \left(\sum_{k=1}^{m_\ell} \phi(\gamma_k^{(\ell)}) \mathbf{v}_k^{(\ell)} \otimes \mathbf{e}_k^{(\ell)} \right) \quad (19)$$

$$= g \cdot \mathcal{A}_\Gamma^{(\ell)}(\mathbf{h}^{(\ell)}). \quad (20)$$

For $\ell = 0$, according to Step 1, it is easy to check

$$\mathcal{A}_\Gamma^{(0)}(g \cdot \mathbf{h}^{(0)}) = g \cdot \mathcal{A}_\Gamma^{(0)}(\mathbf{h}^{(0)}). \quad (21)$$

Therefore,

$$\begin{aligned} \mathcal{A}_\Gamma(g \cdot \mathbf{h}) &= \bigoplus_{\ell=0}^L \mathcal{A}_\Gamma^{(\ell)}(g \cdot \mathbf{h}^{(\ell)}) \\ &= \bigoplus_{\ell=0}^L g \cdot \mathcal{A}_\Gamma^{(\ell)}(\mathbf{h}^{(\ell)}) \\ &= g \cdot \mathcal{A}_\Gamma(\mathbf{h}). \end{aligned} \quad (22)$$

establishing (9). \square

4 Experiments

In this section, we first introduce the experimental setup. Next, we evaluate our method on several popular datasets and compare it against baseline approaches. Finally, we conduct ablation study and assess data efficiency. Additional experiments, including training efficiency, computational overhead and deviation analysis are provided in the Appendix.

4.1 Experimental Setup

Pretrained model Same as previous work, we adopt the MACE (Batatia et al. 2022) as the backbone to evaluate the effectiveness of our proposed MMEA. The pretrained model employed in this work is MACE-OFF (Kovács et al. 2025). The corresponding hyperparameters are summarized in Table 1. More details about MACE are provided in the Extended version.

Hyperparameter	Value
Correlation order	3
Cutoff radius	5.0
Max order for input	3
Max order for product layer	1
Number of hidden channels	128
Number of interaction layers	2
Number of radial basis functions	8

Table 1: Hyperparameter settings of model.

Datasets Following the experimental setup of MACE, we selected benchmark datasets from the molecular potential prediction domain. More detailed information about these datasets is provided in the Extended version. The target properties for each dataset are energy (E) and forces (F). These datasets include:

Firstly, the revised MD-17 dataset (Christensen and Von Lilienfeld 2020). It consists of ten small organic molecules, for which 100,000 structures were computed at DFT accuracy using a very tight SCF convergence and very dense DFT integration grid. The structures were recomputed from the original MD-17 dataset.

Secondly, the 3BPA dataset (Kovács et al. 2021). It contains DFT train test splits of a flexible drug-like organic molecule sampled from different temperature molecular dynamics trajectories. The model is trained on 500 snapshots sampled at 300 K and tested on three independent test sets for each temperature (300 K, 600 K, and 1200 K). The models can also be tested on the challenging task of computing the energy along dihedral rotations of the molecule.

Thirdly, the acetylacetone (AcAc) dataset (Batatia et al. 2022). It contains trajectories of a small reactive molecule sampled at different temperatures. The task is to train on snapshots sampled at 300 K and test on independent test sets sampled at 300 K and 600 K. Moreover the extrapolation is measured both in temperature and along two internal coordinates of the molecule, the hydrogen transfer path and a partially conjugated double bond rotation, which has a very high barrier for rotation.

Experimental Method We compare the Mean Absolute Error (MAE) / Root Mean Square Error (RMSE) of energy and forces for pretrained models with (i) full-parameter fine-tuning (Full), (ii) ELoRA fine-tuning, and (iii) MMEA fine-tuning (our method). Except for the rMD17 dataset, we also recorded the performance of (iv) trained from scratch (Scratch). All fine-tuning methods in each experiment employ the same pretrained model, dataset, and training hyperparameters. Overall, our method follows the experimental setup of ELoRA. The key distinction is that all equivariant linear layers previously fine-tuned with ELoRA are now fine-tuned using our proposed MMEA method. Tensor Product layers and a small number of scalar networks are consistent with the original ELoRA configuration. Unless otherwise specified, the low-rank components in both methods are set to rank $r = 16$. Each experiment is conducted over three runs.

For the rMD17 dataset, since the performance of few-shot learning models is significantly influenced by data splits, we randomly sampled 50 configurations from the officially provided training set to construct our training subset, ensuring that all methods use the exact same data, thereby enabling fair training. For the remaining datasets, the results for full-parameter and ELoRA fine-tuning are directly taken from previous work (Wang et al. 2025).

The experiments were carried out with the PyTorch framework (Paszke et al. 2019) on a server running Ubuntu 22.04.5 LTS. We utilized a single NVIDIA A100 GPU for training and evaluation. More training details are provided in the Extended version.

4.2 Experimental Results

Results on rMD17 The rMD17 dataset here serves as an effective benchmark for evaluating model performance under few-shot learning scenarios. The results (Table 2) show that our proposed MMEA method consistently outperforms previous baselines. Compared to ELoRA, MMEA achieves an additional average improvement of approximately 6.6% on the energy metric and 8.7% on the force metric. Overall, MMEA compresses the average energy and force MAEs by 6–8%, achieves 10–25% improvements on high-error molecules, and establishes new state-of-the-art results. These findings demonstrate that MMEA can more effectively leverage a model’s existing knowledge to attain high-precision predictions in extremely low-sample scenarios.

Results on 3BPA The results of evaluation on 3BPA Dataset are shown in Table 3. The results show that MMEA yields a modest improvement over ELoRA on the 300 K test, that has the same distribution as the training set. When the same model is evaluated directly at higher temperatures, it consistently maintains lower energy and force errors, which demonstrates MMEA’s ability to generalize and mitigate forgetting of pretrained knowledge. Furthermore, on the dihedral slice test, MMEA reduces both energy and force error significantly, indicating a more accurate reconstruction of the potential energy surface and improved stability and precision in conformer predictions.

rMD17		MAE↓		
Molecule	Metric	Full	ELoRA	MMEA
Aspirin	E	9.7	8.0 (↓ 18%)	7.3 (↓ 25%)
	F	23.9	18.3 (↓ 23%)	16.4 (↓ 31%)
Azobenzene	E	4.6	4.0 (↓ 13%)	3.9 (↓ 15%)
	F	14.8	12.6 (↓ 15%)	11.9 (↓ 20%)
Benzene	E	0.3	0.2 (↓ 33%)	0.2 (↓ 33%)
	F	2.4	1.6 (↓ 33%)	1.4 (↓ 42%)
Ethanol	E	2.9	2.3 (↓ 21%)	2.3 (↓ 21%)
	F	14.4	11.5 (↓ 20%)	10.8 (↓ 25%)
Malonaldehyde	E	6.8	6.6 (↓ 3%)	6.4 (↓ 6%)
	F	25.4	22.3 (↓ 12%)	21.6 (↓ 15%)
Naphthalene	E	1.8	1.4 (↓ 22%)	1.2 (↓ 33%)
	F	8.1	6.0 (↓ 26%)	5.7 (↓ 30%)
Paracetamol	E	6.5	4.7 (↓ 28%)	4.3 (↓ 34%)
	F	20.3	14.5 (↓ 29%)	13.3 (↓ 34%)
Salicylic	E	4.3	3.2 (↓ 26%)	2.9 (↓ 33%)
	F	17.2	13.3 (↓ 23%)	12.0 (↓ 30%)
Toluene	E	1.8	1.4 (↓ 22%)	1.2 (↓ 33%)
	F	8.8	6.2 (↓ 30%)	5.4 (↓ 39%)
Uracil	E	2.9	2.1 (↓ 28%)	2.0 (↓ 31%)
	F	15.8	12.3 (↓ 22%)	10.7 (↓ 32%)

Table 2: Experimental results on 10 Molecules of rMD17. Each model is trained on 50 samples. Results are rounded to one decimal place, the decrease ratio is calculated based on the results of full fine-tuning, and the best results are in bold.

Results on AcAc The results of evaluation on AcAc Dataset are shown in Table 4. At 300 K, MMEA achieves the lowest energy and force RMSE, slightly outperforming other baselines. When generalizing to 600 K, it still maintains low errors, showing that it effectively leverages the pre-trained model’s knowledge for better generalization. We also evaluated MMEA with rank set to 32 on this dataset. By increasing its parameter count, we observed a slight performance improvement.

The experiments across multiple datasets reveal the effectiveness of our method. On the rMD17 benchmark, we demonstrate its few-shot learning capability under limited data. On the 3BPA and AcAc benchmarks, we demonstrate its superior generalization performance. This outcome is intuitive: by preserving the multiplicity basis, our method more fully exploits the knowledge in the pretrained model.

4.3 Parameter Efficiency

We compare the parameter efficiency of Full fine-tuning, ELoRA, and our proposed MMEA with two different rank settings. The result is shown in Table 5. Our method ($r = 16$) updates approximately 20% of the parameters used in full fine-tuning, corresponding to about 85% of ELoRA’s parameter budget, while achieving superior fine-tuning performance. When r is set to 32, the parameter count of our method exceeds ELoRA, but as shown in Table 4, it can fur-

3BPA		RMSE↓			
Condition	Metric	Scratch	Full	ELoRA	MMEA
300 K	E	3.0 ± 0.2	3.3 ± 0.03	3.0 ± 0.05	2.7 ± 0.03
	F	8.8 ± 0.3	7.8 ± 0.01	7.5 ± 0.05	7.5 ± 0.01
600 K	E	9.7 ± 0.5	7.3 ± 0.04	6.5 ± 0.10	6.5 ± 0.03
	F	21.8 ± 0.6	16.6 ± 0.05	15.5 ± 0.12	15.4 ± 0.14
1200 K	E	29.8 ± 1.0	20.3 ± 0.17	17.6 ± 0.11	17.1 ± 0.14
	F	62.0 ± 0.7	48.7 ± 0.56	42.0 ± 0.51	39.7 ± 0.12
Dihedral Slices	E	7.8 ± 0.6	7.3 ± 0.28	5.9 ± 0.28	5.6 ± 0.20
	F	16.5 ± 1.7	12.3 ± 0.10	11.4 ± 0.17	10.6 ± 0.17

Table 3: Experimental results under three temperature conditions for 3BPA. The model was trained at 300 K and tested at 300 K, 600 K, and 1200 K to evaluate its generalization performance. Standard deviations are computed over three runs. The best results are in bold.

AcAc		RMSE↓				
Condition	Metric	Scratch	Full	ELoRA	MMEA($r = 16$)	MMEA($r = 32$)
300 K	E	0.9 ± 0.03	1.0 ± 0.02	0.8 ± 0.03	0.7 ± 0.04	0.7 ± 0.02
	F	5.1 ± 0.10	5.1 ± 0.07	4.5 ± 0.06	<u>4.4 ± 0.02</u>	4.2 ± 0.02
600 K	E	4.6 ± 0.3	5.8 ± 0.28	3.9 ± 0.33	<u>3.6 ± 0.06</u>	3.2 ± 0.10
	F	22.4 ± 0.9	16.4 ± 0.70	13.6 ± 0.26	<u>13.2 ± 0.11</u>	13.0 ± 0.23

Table 4: Experimental results under two temperature conditions for AcAc. The model was trained at 300 K and tested at 300 K and 600 K to evaluate its generalization performance. Standard deviations are computed over three runs. The best results are in bold and the second best are underlined.

ther improve performance.

Method	r	# params	% of Full FT
Full	/	751896	100.0%
ELoRA	16	175880	23.4%
MMEA	16	151354	20.1%
MMEA	32	201258	26.7%

Table 5: Trainable parameter for fine-tuning on rMD17.

4.4 Ablation Study

Table 6 presents the ablation study conducted on the rMD17-Aspirin dataset. We removed several key components individually to examine their effects on model performance. w/o nonlinear activation means removing the activation function; w/o input-head reuse indicates disabling the input-head reuse mechanism, where scalar outputs ($l = 0$) and higher-order tensor outputs ($l > 0$) are processed using separate MLPs; w/o high-order refers to omitting the high-order adjustments, i.e., tuning only the scalar information; w/o scalar modulation means no modulation is applied to the scalar channels, which significantly degrades performance; and use same modulation applies identical modulation to all higher-order channels. These results collectively demonstrate that each proposed component plays an important role in improving model performance.

Additionally, we compared two other common fine-tuning strategies. The first (Readout) tunes only the readout (i.e., the layer directly producing the final outputs), updating just 0.3% of the model’s parameters. While extremely parameter-efficient, this approach greatly limits the model’s learning capacity and yields poor fine-tuning performance. The second (Adapter) employs a conventional bottleneck for fine-tuning; by mixing features from different irreducible representations, it breaks the model’s equivariance. As a result, the model cannot fully leverage its pre-trained knowledge and again suffers degraded performance. Both of these methods underperform full fine-tuning, ELoRA, and MMEA.

Like many PEFT approaches, the number of trainable parameters in MMEA is governed by the rank of its bottleneck structure. To examine how this hyperparameter influences model performance, we evaluate forces prediction on the rMD17-Aspirin dataset. As illustrated in Fig. 2, increasing the rank initially improves accuracy, but beyond a certain point leads to performance degradation. This trend reflects a common trade-off: low-rank configurations may lack expressiveness and underfit, whereas higher ranks risk overfitting due to excess capacity. To enable direct comparison with prior work, we use $r = 16$ in our main experiments, to achieve stronger results with fewer parameters. Nevertheless, we observe that a larger setting, $r = 32$, yields better accuracy in this task, suggesting that further improvements

rMD17-Aspirin		MAE↓	
Method	Energy	Forces	
MMEA	7.3	16.4	
Full	9.7	23.9	
w/o nonlinear activation	7.6	16.4	
w/o input-head reuse	9.2	16.7	
w/o scalar modulation	12.9	30.5	
w/o high-order modulation	8.3	16.6	
shared high-order modulation	7.6	16.6	
Readout	23.8	36.8	
Adapter	11.0	26.3	

Table 6: Ablation study results on rMD17-Aspirin Dataset.

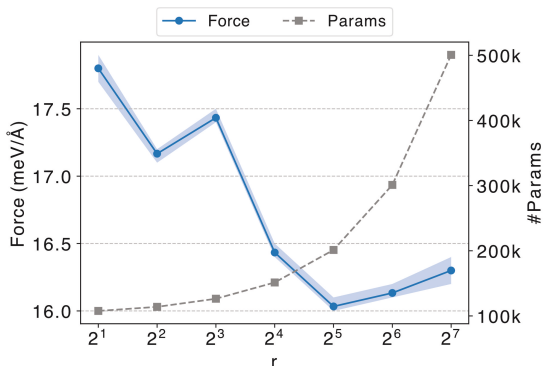


Figure 2: Force MAE and parameter counts across different ranks. Results are reported on the rMD17-Aspirin dataset.

are possible when the parameter budget allows. Therefore, the optimal rank should be treated as a task-specific hyperparameter and tuned accordingly for each downstream application.

5 Conclusion

Equivariant pretrained models are currently indispensable for machine-learning interatomic potentials as they preserve the underlying physical symmetries needed for accurate, reliable atomistic simulations. Fine-tuning these models improves downstream accuracy, especially in low-data scenarios, but standard PEFT techniques break equivariance and thus underperform. Approaches such as ELoRA address this by explicitly retaining symmetry during adaptation, yet their parameter efficiency and performance can still be improved.

In a well-trained equivariant GNN, every irreducible order and its associated multiplicity channels already form a robust representation; effective adaptation therefore requires only subtle, per-channel scaling. Leveraging this insight, we introduce MMEA, a PEFT method that uses scalar gating to modulate each order-multiplicity channel independently, avoiding any mixing within an irreducible order. This preserves equivariance, exploits the pretrained knowledge more

directly, and attains stronger performance with even fewer trainable parameters. MMEA delivers state-of-the-art results on the rMD17, 3BPA, and AcAc benchmarks, demonstrating its effectiveness and efficiency.

Limitations Although the MMEA method demonstrates superior performance on various molecular datasets, we observed certain limitations under specific conditions. Firstly, when the target system significantly deviates from the pretrained model distribution, our method underperforms compared to ELoRA. This occurs primarily because ELoRA permits the mixing of multiplicities, offering greater flexibility and making it easier to adaptively fit challenging datasets. However, for entirely novel systems with minimal overlap to the pretrained dataset, full-parameter fine-tuning remains the most effective strategy, as shown in Table 7. Therefore, and quite intuitively, when the distribution gap between the pre-training model and the downstream task is large, one should train from scratch or perform full fine-tuning; when the gap is moderate or small, leveraging ELoRA or MMEA can yield better performance. Secondly, unlike ELoRA, whose learned weights can be seamlessly merged back into the backbone model parameters, our approach does not allow parameter merging. Consequently, our method introduces slight additional computational overhead during inference.

rMD17		MAE↓		
Molecule	Metric	Full	ELoRA	MMEA
Aspirin	E	11.4	12.5	13.3
	F	29.4	31.1	31.9
Uracil	E	3.4	3.5	3.6
	F	19.7	20.7	21.4

Table 7: Experimental results using MACE-MP (Batatia et al. 2023) pretrained on inorganic compounds to predict organic compounds, assessing extreme out-of-distribution performance. The best results are in bold.

Future Work Future research directions include the following aspects: Firstly, developing efficient equivariant methods for Fully Connected Tensor Product. Since the tensor product involves multiple inputs, how to implement gating effectively remains an open question. We have attempted to apply gating based on the output, but this approach has not shown significant benefits. Secondly, designing a generalizable fine-tuning method. A promising direction is to combine the strengths of the three existing approaches to create a unified method that consistently demonstrates strong performance on data distributions that differ from the pretraining distribution. Thirdly, improving fine-tuning efficiency. Currently, fine-tuning convergence remains relatively slow. Future work could explore the development of more efficient fine-tuning schemes that enable faster convergence.

Acknowledgments

This research was supported by the Research Grants Council of Hong Kong (Grant No. T42-513/24-R).

References

- Anderson, B.; Hy, T. S.; and Kondor, R. 2019. Cormorant: Covariant molecular neural networks. *Advances in neural information processing systems*, 32.
- Aykent, S.; and Xia, T. 2025. Gotennet: Rethinking efficient 3d equivariant graph neural networks. In *The Thirteenth International Conference on Learning Representations*.
- Batatia, I.; Benner, P.; Chiang, Y.; Elena, A. M.; Kovács, D. P.; Riebesell, J.; Advincula, X. R.; Asta, M.; Avaylon, M.; Baldwin, W. J.; et al. 2023. A foundation model for atomistic materials chemistry. *arXiv preprint arXiv:2401.00096*.
- Batatia, I.; Kovacs, D. P.; Simm, G. N. C.; Ortner, C.; and Csanyi, G. 2022. MACE: Higher Order Equivariant Message Passing Neural Networks for Fast and Accurate Force Fields. In Oh, A. H.; Agarwal, A.; Belgrave, D.; and Cho, K., eds., *Advances in Neural Information Processing Systems*.
- Batzner, S.; Musaelian, A.; Sun, L.; Geiger, M.; Mailoa, J. P.; Kornbluth, M.; Molinari, N.; Smidt, T. E.; and Kozinsky, B. 2022. E (3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature communications*, 13(1): 2453.
- Behler, J. 2017. First principles neural network potentials for reactive simulations of large molecular and condensed systems. *Angewandte Chemie International Edition*, 56(42): 12828–12840.
- Behler, J.; and Parrinello, M. 2007. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Physical review letters*, 98(14): 146401.
- Bethune, L.; Grangier, D.; Busbridge, D.; Gualdoni, E.; Curi, M.; and Ablin, P. 2025. Scaling laws for forgetting during finetuning with pretraining data injection. *arXiv preprint arXiv:2502.06042*.
- Burke, K. 2012. Perspective on density functional theory. *The Journal of chemical physics*, 136(15).
- Cen, J.; Li, A.; Lin, N.; Ren, Y.; Wang, Z.; and Huang, W. 2024. Are high-degree representations really unnecessary in equivariant graph neural networks? *Advances in Neural Information Processing Systems*, 37: 26238–26266.
- Chang, J.; and Zhu, S. 2025. MGNN: Moment Graph Neural Network for Universal Molecular Potentials. *npj Computational Materials*, 11(1): 55.
- Christensen, A. S.; and Von Lilienfeld, O. A. 2020. On the role of gradients for machine learning of molecular energies and forces. *Machine Learning: Science and Technology*, 1(4): 045018.
- Devi, R.; Butler, K. T.; and Sai Gautam, G. 2024. Optimal pre-train/fine-tune strategies for accurate material property predictions. *npj Computational Materials*, 10(1): 300.
- Elfwing, S.; Uchibe, E.; and Doya, K. 2018. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural networks*, 107: 3–11.
- Fiedler, L.; Modine, N. A.; Schmerler, S.; Vogel, D. J.; Popoola, G. A.; Thompson, A. P.; Rajamanickam, S.; and Cang, A. 2023. Predicting electronic structures at any length scale with machine learning. *npj Computational Materials*, 9(1): 115.
- Gasteiger, J.; Becker, F.; and Günnemann, S. 2021. Gemnet: Universal directional graph neural networks for molecules. *Advances in Neural Information Processing Systems*, 34: 6790–6802.
- Gasteiger, J.; Groß, J.; and Günnemann, S. 2020. Directional message passing for molecular graphs. *arXiv preprint arXiv:2003.03123*.
- Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; and Dahl, G. E. 2017. Neural message passing for quantum chemistry. In *International conference on machine learning*, 1263–1272. Pmlr.
- Houlsby, N.; Giurgiu, A.; Jastrzebski, S.; Morrone, B.; De Laroussilhe, Q.; Gesmundo, A.; Attariyan, M.; and Gelly, S. 2019. Parameter-efficient transfer learning for NLP. In *International conference on machine learning*, 2790–2799. PMLR.
- Howard, J.; and Ruder, S. 2018. Universal language model fine-tuning for text classification. *arXiv preprint arXiv:1801.06146*.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Kovács, D. P.; Moore, J. H.; Browning, N. J.; Batatia, I.; Horton, J. T.; Pu, Y.; Kapil, V.; Witt, W. C.; Magdau, I.-B.; Cole, D. J.; et al. 2025. Mace-off: Short-range transferable machine learning force fields for organic molecules. *Journal of the American Chemical Society*, 147(21): 17598–17611.
- Kovács, D. P.; Oord, C. v. d.; Kucera, J.; Allen, A. E.; Cole, D. J.; Ortner, C.; and Csányi, G. 2021. Linear atomic cluster expansion force fields for organic molecules: beyond rmse. *Journal of chemical theory and computation*, 17(12): 7696–7711.
- Lee, B. K.; Mayhew, E. J.; Sanchez-Lengeling, B.; Wei, J. N.; Qian, W. W.; Little, K. A.; Andres, M.; Nguyen, B. B.; Moloy, T.; Yasonik, J.; et al. 2023. A principal odor map unifies diverse tasks in olfactory perception. *Science*, 381(6661): 999–1006.
- Li, S.; Han, X.; and Bai, J. 2024. Adapterggn: Parameter-efficient fine-tuning improves generalization in gnns. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 13600–13608.
- Liao, Y.-L.; and Smidt, T. 2022. Equiformer: Equivariant graph attention transformer for 3d atomistic graphs. *arXiv preprint arXiv:2206.11990*.
- Liu, Y.; Zeng, K.; Wang, H.; Song, X.; and Zhou, B. 2021. Content matters: A GNN-based model combined with text semantics for social network cascade prediction. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 728–740. Springer.
- Maruf, M. U.; Kim, S.; and Ahmad, Z. 2025. Equivariant Machine Learning Interatomic Potentials with Global Charge Redistribution. *arXiv preprint arXiv:2503.17949*.

- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Perez, E.; Strub, F.; De Vries, H.; Dumoulin, V.; and Courville, A. 2018. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Satorras, V. G.; Hoogeboom, E.; and Welling, M. 2021. E(n) equivariant graph neural networks. In *International conference on machine learning*, 9323–9332. PMLR.
- Schütt, K.; Unke, O.; and Gastegger, M. 2021. Equivariant message passing for the prediction of tensorial properties and molecular spectra. In *International conference on machine learning*, 9377–9388. PMLR.
- Schütt, K. T.; Sauceda, H. E.; Kindermans, P.-J.; Tkatchenko, A.; and Müller, K.-R. 2018. SchNet—a deep learning architecture for molecules and materials. *The Journal of chemical physics*, 148(24).
- Smith, J. S.; Isayev, O.; and Roitberg, A. E. 2017. ANI-1: an extensible neural network potential with DFT accuracy at force field computational cost. *Chemical science*, 8(4): 3192–3203.
- Thölke, P.; and De Fabritiis, G. 2022. Torchmd-net: equivariant transformers for neural network based molecular potentials. *arXiv preprint arXiv:2202.02541*.
- Thomas, N.; Smidt, T.; Kearnes, S.; Yang, L.; Li, L.; Kohlhoff, K.; and Riley, P. 2018. Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds. *arXiv preprint arXiv:1802.08219*.
- Unke, O. T.; Chmiela, S.; Gastegger, M.; Schütt, K. T.; Sauceda, H. E.; and Müller, K.-R. 2021. SpookyNet: Learning force fields with electronic degrees of freedom and non-local effects. *Nature communications*, 12(1): 7273.
- Unke, O. T.; and Meuwly, M. 2019. PhysNet: A neural network for predicting energies, forces, dipole moments, and partial charges. *Journal of chemical theory and computation*, 15(6): 3678–3693.
- Wang, C.; Hu, S.; Tan, G.; and Jia, W. 2025. ELoRA: Low-Rank Adaptation for Equivariant GNNs. In *International Conference on Machine Learning*.
- Wieder, O.; Kohlbacher, S.; Kuenemann, M.; Garon, A.; Ducrot, P.; Seidel, T.; and Langer, T. 2020. A compact review of molecular property prediction with graph neural networks. *Drug Discovery Today: Technologies*, 37: 1–12.
- Wu, S.; Sun, F.; Zhang, W.; Xie, X.; and Cui, B. 2022. Graph neural networks in recommender systems: a survey. *ACM Computing Surveys*, 55(5): 1–37.
- Zaken, E. B.; Ravfogel, S.; and Goldberg, Y. 2021. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. *arXiv preprint arXiv:2106.10199*.
- Zhang, L.; Han, J.; Wang, H.; Car, R.; and E, W. 2018. Deep potential molecular dynamics: a scalable model with the accuracy of quantum mechanics. *Physical review letters*, 120(14): 143001.
- Zhou, G.; Gao, Z.; Ding, Q.; Zheng, H.; Xu, H.; Wei, Z.; Zhang, L.; and Ke, G. 2023. Uni-Mol: A Universal 3D Molecular Representation Learning Framework. In *The Eleventh International Conference on Learning Representations*.