

Towards Faster and More Compact Foundation Models for Molecular Property Prediction

Yasir Ghunaim, Andrés Villa, Gergo Ignacz, Gyorgy Szekely, Motasem Alfarra, and Bernard Ghanem
King Abdullah University of Science and Technology (KAUST)

Abstract

Advancements in machine learning for molecular property prediction have improved accuracy but at the expense of higher computational cost and longer training times. Recently, the Joint Multi-domain Pre-training (JMP) foundation model has demonstrated strong performance across various downstream tasks with reduced training time over previous models. Despite JMP’s advantages, fine-tuning it on molecular datasets ranging from small-scale to large-scale requires considerable time and computational resources. In this work, we investigate strategies to enhance efficiency by reducing model size while preserving performance. To better understand the model’s efficiency, we analyze the layer contributions of JMP and find that later interaction blocks provide diminishing returns, suggesting an opportunity for model compression. We explore block reduction strategies by pruning the pre-trained model and evaluating its impact on efficiency and accuracy during fine-tuning. Our analysis reveals that removing two interaction blocks results in a minimal performance drop, reducing the model size by 32% while increasing inference throughput by 1.3×. These results suggest that JMP-L is over-parameterized and that a smaller, more efficient variant can achieve comparable performance with lower computational cost. Our study provides insights for developing lighter, faster, and more scalable foundation models for molecular and materials discovery. The code is publicly available at: github.com/Yasir-Ghunaim/efficient-jmp.

1. Introduction

Molecular property prediction using density functional theory (DFT) and molecular dynamics (MD) calculations plays a crucial role in the discovery of novel materials, including pharmaceutical drugs [32], catalysts [8, 28, 44], metal-organic frameworks [31], and polymers [37]. However, the high computational cost of DFT and MD calculations limits their feasibility for large-scale, high-throughput

searches. To address this challenge, machine learning potentials have been developed to accelerate DFT and MD calculations [4, 6], leveraging the latest large-scale datasets [23], such as Open Catalyst 2020 (OC20) [8], Open Catalyst 2022 (OC22) [44], and ODAC23 [41]. However, training models from scratch for different tasks remains a major bottleneck for their widespread adoption. Differences in applied DFT theories, molecular system sizes, and chemical diversity increase the molecular system complexity, thus hindering the generalizability and scalability of machine learning models in chemistry.

Recent progress in efficient pre-training strategies [47, 50], the availability of extensive DFT and MD datasets [8, 15, 40, 44], and the introduction of specialized chemical benchmarks [14, 33] have led to the emergence of foundation models for molecular property prediction. Foundation models such as the Joint Multi-domain Pre-training (JMP) model [38] and MACE-MP-0 [1–3] have demonstrated outstanding performance in diverse molecular tasks. In particular, JMP adapts the pre-train-then-finetune paradigm from vision and language tasks to molecular property prediction. By pre-training on large datasets, the JMP model captures generalizable molecular representations that enable faster fine-tuning for downstream tasks, overcoming the need to train models from scratch for each new application.

Although the large variant of JMP (JMP-L) has outperformed state-of-the-art models on 34 out of 40 tasks, its efficiency in fine-tuning and inference has yet to be addressed. With 160M parameters, JMP-L achieves similar performance to MACE [2], which uses only 3M parameters, suggesting potential over-parameterization. This over-parameterization increases memory and compute requirements and leads to higher carbon emissions [38], reducing overall sustainability. Although 160M parameters are relatively small compared to vision and language models, the parameter-to-data ratio in molecular ML remains disproportionately large. For instance, MD17 contains only 1,000 training samples with an average of 13 nodes per graph [38], making a model of this size inefficient for small datasets.

To address these limitations, we perform an in-depth analysis of the efficiency of JMP-L. By examining its in-

Correspondence to: yasir.ghunaim@kaust.edu.sa

teraction block hierarchy, we find that higher-order blocks contribute less to overall performance. This observation aligns with recent findings in large language models, where deeper layers often yield diminishing returns [19]. This motivates our exploration of block reduction, a pruning strategy that removes the least important layers to improve efficiency while maintaining accuracy. Additionally, we investigate knowledge distillation techniques tailored to molecular graph neural networks, integrating them with block reduction to assess their combined impact. Although pruning [26] and distillation [48] are widely used in other domains, their application to molecular property prediction, particularly within the pre-train-then-finetune paradigm, remains underexplored.

Our findings reveal that pruning and distillation improve JMP-L’s efficiency while preserving comparable performance for most tasks. Specifically, we show that a pruned and distilled variant of JMP-L achieves comparable accuracy to the original model across in-distribution and out-of-distribution downstream tasks. By removing two interaction blocks, we reduce the model size by 32%, decreasing the parameter count from 160M to 108M, while improving inference throughput by 1.3× compared to the baseline model. These results confirm that JMP-L is over-parameterized for many tasks, and smaller, more efficient versions can achieve similar performance with reduced computational cost. In summary, our contributions are three-fold:

- We reduce the number of JMP-L parameters by 32% to 108M and achieve 1.3× faster inference while maintaining performance.
- We evaluate the impact of block reduction and knowledge distillation on pre-training across in-distribution and out-of-distribution downstream tasks.
- We demonstrate that later interaction blocks of JMP-L contribute less to performance, supporting the case for model compression.

2. Related Work

2.1. Foundation Models in Molecular Property Prediction

Pre-trained models have made considerable advancements in developing robust architectures across various domains. Notable examples in the vision domain include ResNet [21] and ViT [13], which leverage large-scale datasets such as ImageNet [12] to enhance image processing. In contrast, deep learning models for molecular property prediction have primarily been task-specific [3, 24], limiting their utility as general-purpose pre-trained models. Recently, JMP [38] introduced a supervised pre-training strategy on large datasets, establishing a shared knowledge base for various downstream tasks. Built on GemNet-OC [18], JMP is the first large-scale foundation model for molecular property

prediction. However, its fine-tuning efficiency remains a challenge, as it requires more than 275 GPU hours to converge [38]. In this work, we provide a comprehensive analysis of JMP and propose a more efficient approach to reduce its computational demands, enhancing its accessibility and scalability for broader applications.

2.2. Efficient Training

2.2.1. Pruning

Pruning is a technique used to reduce the size and complexity of a neural network by eliminating weights, neurons, layers, or filters without compromising accuracy [5, 9, 39]. It is particularly effective when a model is over-parameterized for its task [39]. Structured pruning, which removes entire layers or filters, has been shown to improve memory and computational efficiency in various architectures, including large language models [42, 49], vision transformers [46], and graph neural networks (GNNs) [26]. JMP-L, which is based on the GemNet-OC architecture [18], consists of an embedding layer, six interaction layers, and three MLP layers. Drawing inspiration from pruning techniques in other domains, we investigate the impact of removing GemNet-OC interaction layers to accelerate fine-tuning and inference while maintaining model performance.

2.2.2. Knowledge distillation

Knowledge distillation (KD) is a model compression technique that transfers knowledge from a larger teacher model to a smaller student model, aiming to achieve similar performance with reduced computational costs [7, 22]. Initially introduced by Bucilua *et al.* [7] and later popularized by Hinton *et al.* [22], KD has been widely applied in the language [45], vision [20], and general graph domains [43], mainly in classification tasks. However, applying KD to large-scale regression tasks such as DFT and MD simulations presents unique challenges [16]. Molecular GNNs operate on structured graph data with features distributed across nodes and edges, making direct knowledge transfer from teacher to student challenging [16]. These challenges are amplified when the teacher and student models differ considerably in architecture, making feature alignment more difficult. To address these challenges, Ekström *et al.* [16] propose specialized loss functions—*node2node* ($n2n$), *edge2edge* ($e2e$), *edge2node* ($e2n$), and *vector2vector* ($v2v$)—to supplement standard loss functions and enhance the effectiveness of KD in molecular GNNs. These strategies help bridge the gap between teacher and student models, improving knowledge transfer in complex molecular systems.

The standard loss function \mathcal{L}_0 for molecular GNNs, as outlined in Eq. 1, accounts for both energy and force predictions:

$$\mathcal{L}_0 = \alpha_E \mathcal{L}_E(\hat{E}, E) + \alpha_F \mathcal{L}_F(\hat{F}, F) \quad (1)$$

where E and F represent the ground-truth energy and forces, while \hat{E} and \hat{F} denote their predicted counterparts. The terms \mathcal{L}_E and \mathcal{L}_F denote the energy and force loss functions, respectively, weighted by $\alpha_E, \alpha_F \in \mathbb{R}$.

For knowledge distillation [22], the loss function in Eq. 1 is augmented with an auxiliary distillation loss \mathcal{L}_{KD} , resulting in the following formulation:

$$\mathcal{L} = \mathcal{L}_0 + \lambda \mathcal{L}_{\text{KD}}.$$

In this work, we aim to develop a more efficient variant of the foundational model JMP-L without compromising its performance. Since JMP is built on GemNet-OC [18], we focus on the *node2node* and *edge2edge* losses as key components of the distillation process. Unlike previous studies that apply distillation only at the final task level, our primary objective is to enhance the efficiency of the foundational model itself while systematically assessing its impact on downstream tasks. Specifically, we examine how distillation influences the model’s generalization capabilities, providing deeper insights into its performance across diverse molecular property prediction tasks.

2.2.3. Pruning Coupled with Distillation

Aggressive structured pruning can substantially degrade model performance. For instance, brute-force structural pruning methods, such as L2-based filter-wise pruning, have led to a 50-fold performance drop in LLMPrunner [27]. However, in large models, aggressive pruning combined with fine-tuning can considerably reduce the number of layers—sometimes by half—while incurring only minimal performance loss [19]. Techniques such as parameter-efficient fine-tuning, quantization, and low-rank adapters further help preserve model accuracy post-pruning [19]. Recent KD approaches, such as those proposed by Ekström *et al.* [16], require training the student model from scratch, demanding substantial computational resources. To the best of our knowledge, no prior work has explored the combined use of pruning and distillation for DFT and MD molecular property prediction. Our approach applies distillation to a pre-trained, block-reduced network, offering the potential for improved accuracy while greatly reducing both training and inference time.

3. Block Reduction for Efficient Foundation Models

3.1. Preliminaries

Our work is based on the GemNet-OC architecture [18], although our analysis can be applied to similar architectures.

In particular, we define $f_\theta : \mathbb{R}^{4 \times n} \rightarrow \mathbb{R}^{n \times d}$ as a function that maps a molecular graph, represented by the 3D positions and atomic numbers of n atoms, to a feature space. The feature extraction process is formulated as follows:

$$f(x) = \text{concat}(f_1(x), f_2 \circ f_1(x), \dots, f_b \circ f_{b-1} \circ \dots \circ f_2 \circ f_1(x)) \quad (2)$$

where the model extracts features through b sequential blocks. Each block $f_i : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}^{n \times d}$ (for $i > 1$) refines the representations, while the initial block $f_1 : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times d}$ is an embedding layer that performs the initial transformation. The resulting feature space is of dimension $\mathbb{R}^{n \times (d \times b)}$, obtained by concatenating outputs from all b blocks. This extracted feature representation is then processed through a sequence of multilayer perceptron (MLP) layers, known as FinalMLP in GemNet-OC.

$$g(x) = g_m \circ g_{m-1} \circ \dots \circ g_2 \circ g_1(x)$$

where $g_1 : \mathbb{R}^{n \times (d \times b)} \rightarrow \mathbb{R}^{n \times d}$ transforms the concatenated features into d , with $g(x) : \mathbb{R}^{n \times (d \times b)} \rightarrow \mathbb{R}^{n \times d}$ providing the final transformation. Finally, the output of $g(x)$ is passed to a prediction head $h(x)$, which predicts the three-dimensional force vector for each atom and the molecule’s energy. The full model is the composition between the feature extractor and the MLP layers given by the following:

$$F(x) = h \circ g \circ f(x). \quad (3)$$

For the GemNet-OC architecture used by JMP-L, $b = 7$ (one embedding layer and six interaction blocks) and $m = 5$, resulting in 160.1M parameters in total. JMP-L is shown on top in Figure 1.

3.2. Interaction Block Importance

Interaction blocks are a fundamental component of machine learning potential models (e.g., SchNet [35], GemNet-OC [17]), enabling richer representations and capturing long-range atomic interactions. These models stack multiple interaction blocks in a sequential manner to build higher-body representations and model complex atomic relationships effectively. However, quantifying each interaction block’s contribution to the final prediction is not straightforward, as interactions are highly interdependent and difficult to isolate.

To address this, we propose an approach to measure interaction block importance within the GemNet-OC backbone used by JMP-L. We employ GradCAM [36] to assess each block’s impact on the final output. First, we extract and concatenate the output features from all b blocks, as formulated in Eq. 2. Using these features (f) and Eq. 3, we compute the model’s output and its corresponding loss, \mathcal{L}_0 .

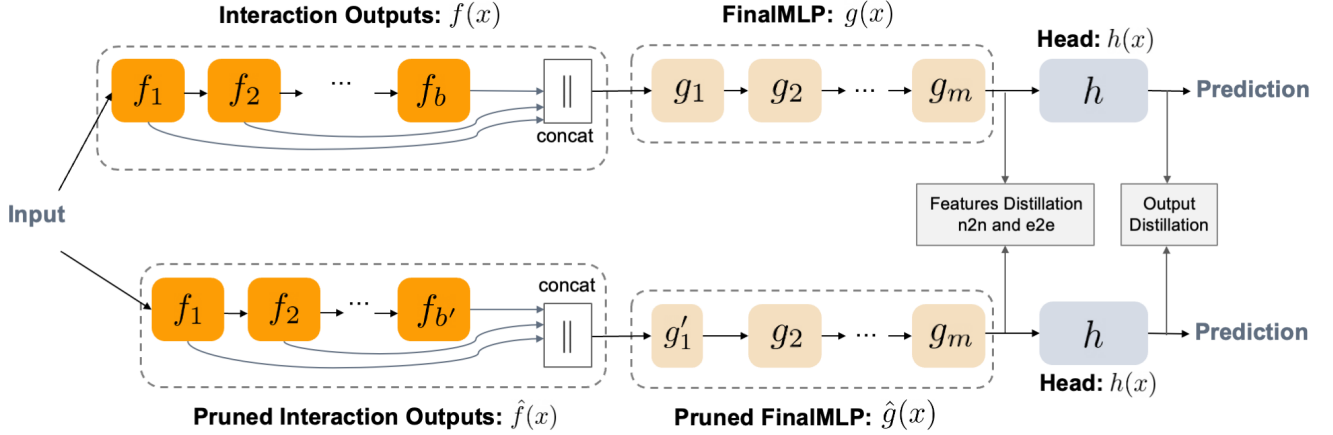


Figure 1. **Block Reduction for Efficient Foundation Models.** The top model represents the foundation model JMP-L, where interaction blocks extract features, which are concatenated and processed by FinalMLP before making predictions. The bottom model is its pruned version, constructed by removing low-importance blocks and adjusting FinalMLP. To mitigate performance degradation, we apply both feature distillation (node-to-node and edge-to-edge) and output distillation to transfer knowledge from the original model.

We then compute the gradient ∇_{CAM} of f with respect to \mathcal{L}_0 and determine each block’s relevance r using:

$$r = \text{ReLU}(f \circ \nabla_{\text{CAM}}).$$

Following the GradCAM methodology, we apply a ReLU activation to emphasize features that positively contribute to the model’s prediction. To quantify the contribution of each interaction block, we decompose the relevance map r , which represents the overall importance of features, into b partitions. Each partition r_i corresponds to a feature dimension d of a specific interaction block. Finally, we compute the importance score for each block by averaging r_i across both the feature and batch dimensions, providing a measure of its overall contribution to the final prediction.

After using GradCAM, we normalize the contributions by dividing each individual contribution by the sum of all contributions. This ensures that the importance scores are comparable across different models and configurations.

3.3. Block Reduction Strategies

Although the pre-trained GemNet-OC (JMP-L) achieves strong performance when fine-tuned across various tasks and datasets, fine-tuning it remains computationally expensive. This inefficiency arises from its large architecture and the high computational cost of each forward pass. For instance, fine-tuning JMP-L on rMD17—containing only 1,000 graphs—still requires 160.1M parameters. Given the small scale of rMD17, this parameter count appears disproportionately large, limiting the practicality of leveraging such a powerful pre-trained model efficiently.

In this work, we aim to improve the efficiency of foundation pre-trained models through block reduction. Specif-

ically, we explore different strategies to construct a reduced model $\hat{F}(x) = h \circ \hat{g} \circ \hat{f}(x)$, where:

$$\hat{f}(x) = \text{concat}(f_1(x), f_2 \circ f_1(x), \dots, f_{b'} \circ f_{b'-1} \circ \dots \circ f_2 \circ f_1(x))$$

and

$$\hat{g}(x) = g_m \circ g_{m-1} \circ \dots \circ g_2 \circ g'_1(x)$$

with $g'_1 : \mathbb{R}^{n \times (d \times b')} \rightarrow \mathbb{R}^{n \times d}$ and $b' < b$.

This formulation reduces the original architecture by removing the last $b - b'$ interaction blocks and adjusting the dimensionality of the first MLP block (g'_1) accordingly. However, removing interaction blocks disrupts the alignment between the feature extractor \hat{f} and the FinalMLP \hat{g} , as it alters the structure of the extracted features. To address this misalignment, we explore three main strategies to restore compatibility between the reduced feature extractor and the FinalMLP.

Random MLP. Random MLP is the simplest baseline, where we resize the first MLP layer g'_1 and randomly initialize its weights. This approach assumes that the features extracted by the remaining interaction blocks are still useful and that the weights of g'_1 can be learned effectively during fine-tuning. We refer to this strategy as RandomMLP in the experimental section.

Sliced MLP. In the Sliced MLP strategy, we retain the parameters of the first MLP layer g'_1 from the original model, truncating it to match the reduced dimensionality of the features. This assumes that the preserved parameters provide a good initialization for fine-tuning, maintaining continuity

between the pre-trained and pruned model. Unless otherwise stated, all of our block reduction experiments follow this strategy.

Knowledge Distillation. In the knowledge distillation approach, we introduce a learning paradigm tailored for block reduction. Specifically, we follow [16] by distilling the force predictions of the pre-trained model F into its block-reduced counterpart \hat{F} . This is achieved by optimizing the following objective:

$$\min_{\theta} \mathbb{E}_{x \sim \mathcal{D}} \|\hat{F}_{\theta}(x) - F_{\theta_0}(x)\|_1$$

where \mathcal{D} represents the data distribution used during the pre-training phase of F .

To further align the representations between the original and pruned models, we incorporate node-to-node (n2n) and edge-to-edge (e2e) distillation. This extends the objective to:

$$\min_{\theta} \mathbb{E}_{x \sim \mathcal{D}} \left[\|\hat{F}(x; \theta) - F(x; \theta_0)\|_1 + \sum_{i=1}^{n'} \|\hat{g}_i(\hat{f}(x); \theta) - g_i(f(x); \theta_0)\|_1 \right].$$

The first term represents output distillation, while the second ensures feature-level consistency between the pruned and original models. Our full pipeline is illustrated in Figure 1.

It is important to note that all prior block reduction approaches operate on the pre-trained foundational model. Thus, block reduction produces a generalist model, which must still undergo fine-tuning on task-specific datasets to become a specialized model for a given downstream application.

4. Experiments

4.1. Deeper layers contribute less

To develop a more efficient pre-trained version of JMP-L for diverse downstream tasks, we analyze the contribution of each interaction block to the final output prediction on the pre-training distribution, as outlined in Sec.3.2. Figure 2 reveals a gradual decline in relevance in deeper layers, with the sixth and seventh blocks exhibiting the lowest contribution. This gradual decline suggests that blocks with low contribution can be removed with minimal impact on performance. Although the embedding block (f_1) also shows low relevance, it directly feeds into the first interaction block (f_2), which initializes message passing and is crucial for propagating information between atoms and bonds. Removing the embedding block could introduce major structural disruptions, potentially degrading performance, and is therefore not considered for block reduction.

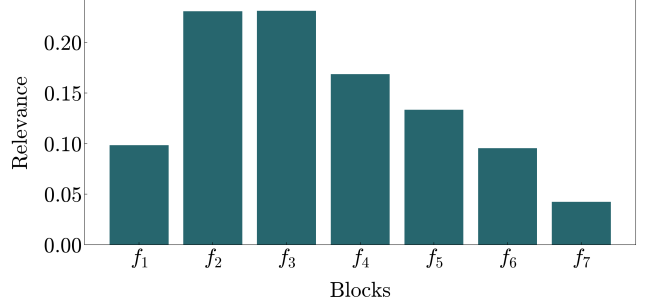


Figure 2. **Block Relevance Analysis.** This table illustrates the contribution of each output block in JMP-L to the final prediction. The first block, f_1 , represents the embedding output, while f_2 to f_7 correspond to the six interaction blocks of GemNet-OC. The results show diminishing returns in deeper interaction blocks, suggesting they are strong candidates for pruning to improve efficiency without considerable performance loss.

4.2. Block Reduction and Distillation in Pre-training

We investigate the effectiveness of block reduction (BR) and knowledge distillation (KD) during the pre-training phase. The importance analysis of the blocks shows that the deeper interaction blocks in JMP-L contribute less to the final prediction compared to the earlier ones. This raises a key question: how much do these later blocks impact prediction performance, and if their removal leads to degradation, to what extent can knowledge distillation recover the lost performance? To address these questions, we explore a combined approach of block reduction and knowledge distillation, assessing whether distilling knowledge from the full model into a block-reduced version can maintain performance while reducing computational costs.

Settings: We use the same pre-training datasets as in JMP [38], including OC20 [8], OC22 [44], ANI-1x [40], and Transition-1x [34], with a total of 120M training samples. Following our block reduction (BR) strategy, we sequentially remove interaction blocks starting from the last one, as indicated by the importance analysis in the previous section. This allows us to construct progressively smaller versions of JMP-L, retaining 5, 4, 3, and 2 interaction blocks. To mitigate potential performance degradation, we apply a brief knowledge distillation (KD) phase to each pruned model using less than 1.5% of the pre-training datasets. We find that running KD for under 2 GPU-days on an A100 is sufficient for convergence. We report performance using the mean absolute error (MAE).

Observations: Table 1 presents the results of block reduction (BR) and knowledge distillation (KD) during the pre-training stage. Removing interaction blocks leads to a performance drop proportional to the number of blocks removed. Notably, OC20 and OC22 show smaller performance drops compared to other datasets, likely due to their

Table 1. **Force MAE Evaluation During Pre-training.** Impact evaluation of our block reduction (BR) approach on pre-training performance, comparing results with and without knowledge distillation (KD). The reported values represent the force MAE (meV/Å) across the OC20, OC22, ANI-1x, and Transition-1x datasets used for pre-training. KD reduces the performance drop relative to the teacher model. Note: The block count excludes the embedding output.

	# of Blocks	OC20	OC22	ANI-1x	Transition-1x
BR	2	106.6 (-89.6)	111.9 (-89.0)	711.8 (-689.3)	188.3 (-175.5)
+ KD	2	52.1 (-35.1)	56.3 (-33.4)	144.4 (-121.9)	53.3 (-40.6)
BR	3	94.1 (-77.1)	99.3 (-76.3)	615.1 (-592.6)	149.6 (-136.9)
+ KD	3	42.1 (-25.2)	45.6 (-22.7)	97.6 (-75.1)	39.1 (-26.4)
BR	4	65.3 (-48.4)	68.9 (-46.0)	443.4 (-421.0)	99.8 (-87.1)
+ KD	4	26.7 (-9.8)	34.5 (-11.6)	58.1 (-35.7)	24.2 (-11.4)
BR	5	39.1 (-22.1)	45.7 (-22.7)	220.0 (-197.6)	48.5 (-35.8)
+ KD	5	19.8 (-2.8)	25.8 (-2.9)	30.1 (-7.7)	16.0 (-3.3)
JMP-L (Teacher)	6	17.0	22.9	22.5	12.7

higher force loss weight during pre-training [38], making the model less sensitive to deeper block removal. Applying both BR and KD greatly reduces the performance gap with the teacher model. For instance, with 5 blocks, the performance difference narrows to just -2.8 and -2.9 meV/Å for OC20 and OC22, respectively. These results demonstrate the effectiveness of combining BR and KD in maintaining predictive accuracy during pre-training. Next, we analyze how well these pruned models perform on downstream tasks.

4.3. Main Results

We evaluate pruned versions of JMP-L across various downstream tasks and knowledge transfer strategies to determine the most efficient fine-tuning approach for optimal performance. A key question is whether distilling during pre-training is more effective than applying block reduction alone. To explore this, we consider different baselines during fine-tuning.

Baselines: We evaluate the following fine-tuning strategies:

- BR (Block Reduction): Remove interaction blocks and slice their corresponding weights in the first layer of FinalMLP.
- BR/RandomMLP: A simpler variant of BR, where instead of pruning the first layer of FinalMLP, we randomly initialize a smaller version to match the reduced number of interaction blocks.
- BR+KD (Block Reduction + Knowledge Distillation): Load the pruned and distilled version of JMP-L, where knowledge distillation has been applied during pre-training.

We further report the original performance of JMP-L from [38] alongside our reproduced version. Our fine-tuning process is constrained by a fixed computational budget, compared to the vanilla JMP-L:

Settings: We evaluate the baselines on a representative set of targets from the datasets used in [38], specifically: Aspirin (rMD17 [10]), U_0 (QM9 [29]), Solvated Amino Acids (SPICE [15]), Ac-Ala3-NHMe (MD22 [11]), and Band Gap (QMOF [30]). For a fair comparison, since the baselines differ in computational demands, we fine-tune each model for 1 GPU-day on a V100, except for QM9, which requires 2 GPU-days to approach convergence. We then evaluate the models on the test set of the corresponding dataset and target. Our results are shown in Figure 3.

Block Reduction is a Strong Baseline: Figure 3 highlights the surprising effectiveness of BR (shown in blue) across datasets. In particular, the 5-block model matches the performance of the 6-block reproduced JMP-L baseline. Even the 4-block model remains competitive across most tasks and outperforms the original JMP-L on the QM9 target. These results suggest that for tasks with longer convergence times, such as QM9, a compressed model not only reduces computational costs but may also converge faster and even surpass the full model’s performance. However, when further reducing to 3 or 2 blocks, we observe a more pronounced drop in performance, indicating that the model may be underfitting the task.

Pre-training Distillation Works in Certain Scenarios: While KD improved performance on the pre-training datasets (as discussed in Section 4.2), its effectiveness on downstream tasks varies, as shown by BR+KD (yellow in

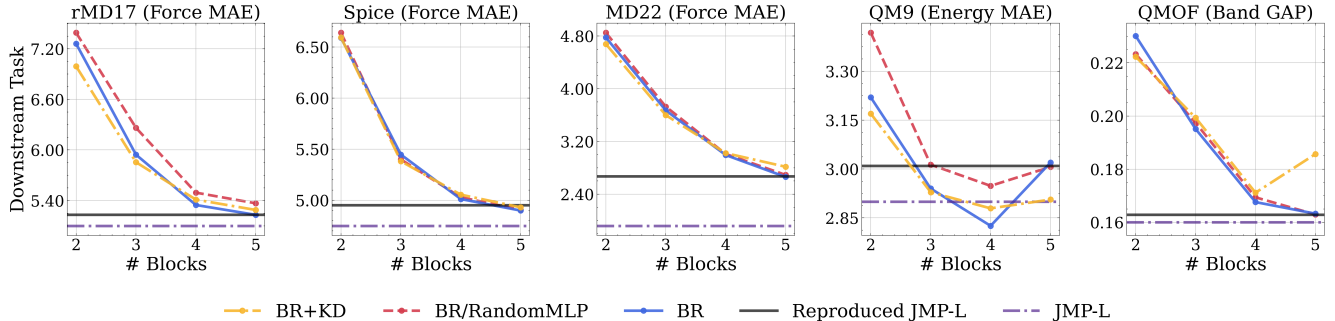


Figure 3. **Evaluation on downstream tasks.** Evaluation of the performance across various downstream tasks using different block reduction strategies: block reduction (BR), block reduction with a randomly initialized MLP (BR/RandomMLP), and block reduction combined with knowledge distillation (BR+KD). Performance is measured in MAE: $\text{meV}/\text{\AA}$ for force targets, meV for the QM9 energy target, and eV for the QMOF band gap target. The original JMP-L model utilizes 6 blocks.

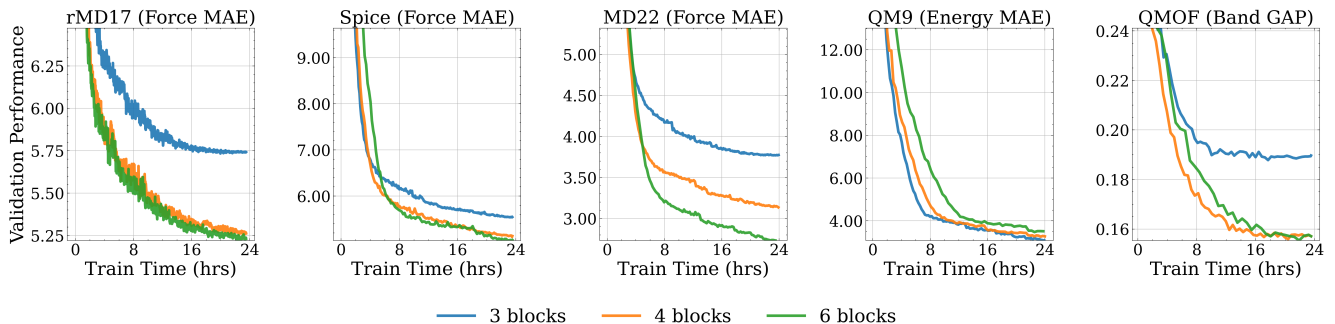


Figure 4. **Training Efficiency Analysis.** We compare the convergence speed of JMP-L models with 3, 4, and 6 blocks over a fixed 1 GPU-day training budget. The 4-block model achieves faster convergence than the 6-block model on QM9 after 8 hours but lags behind on MD22, while performing comparably on other datasets. Performance is measured in MAE: $\text{meV}/\text{\AA}$ for force targets, meV for the QM9 energy target, and eV for the QMOF band gap target.

Figure 3). For example, KD improves performance on rMD17 when using 2 or 3 blocks, but it hurts the performance with 4 and 5 blocks. This could be due to the fact that the last two blocks of JMP-L contribute less to downstream tasks (as indicated by BR results), meaning that distilling from these less relevant blocks during pre-training may introduce noise and distort useful features. Interestingly, KD improves performance for the 5-block model in QM9, suggesting a potential edge case where distillation benefits from specific task characteristics.

JMP-L’s FinalMLP Layer May Indicate Distribution Shifts: The BR/RandomMLP baseline (red in Figure 3) exhibits inconsistent behavior across different tasks. In QMOF, SPICE, and MD22, randomly initializing the first layer of FinalMLP had little impact on performance, suggesting a distribution shift between pre-training and downstream tasks. In contrast, BR/RandomMLP shows a noticeable performance drop in rMD17 and QM9, indicating that the learned FinalMLP features are more relevant to these tasks. This observation also aligns with the improved performance of BR+KD in rMD17 and QM9, where KD effectively preserves useful representations. These findings

suggest that FinalMLP features could serve as indicators of distribution differences across tasks, highlighting variations in task similarity to the pre-training distribution.

4.4. Evaluating Training and Inference Efficiency

Training Efficiency: We complement our analysis with both training and inference speed-ups gains. For training efficiency, Figure 4 illustrates the convergence speed over a fixed budget of 1 GPU-day for models with 3, 4, and 6 blocks across five datasets: rMD17, SPICE, MD22, QM9, and QMOF. The performance is measured in force MAE ($\text{meV}/\text{\AA}$) for rMD17, SPICE, and MD22, energy MAE (meV) for QM9, and band gap error (eV) for QMOF.

In general, reducing the number of interaction blocks accelerates training. The 4-block model has similar convergence behavior to the full 6-block model across most datasets. For example, on QM9, the 4-block model outperforms the 6-block model after 8 hours of training, achieving lower energy MAE. However, on MD22, the 4-block model lags behind, with the 6-block model maintaining a noticeable performance advantage. In contrast, the 3-block model shows the fastest convergence across datasets but exhibits

a worse final validation loss, indicating a trade-off between efficiency and accuracy.

These findings emphasize the importance of selecting an optimal block count to balance convergence speed and final model accuracy. The 4-block model strikes a balance, providing a moderate training speedup over the 6-block model in some datasets while maintaining competitive performance across all datasets.

Table 2. Inference Efficiency Analysis. We evaluate the impact of block reduction on JMP-L’s efficiency using a subset of the QMOF validation set. Reducing interaction blocks lowers computational cost and improves inference speed. The 4-block model provides the best trade-off, achieving a 1.3x speedup with a 32% reduction in parameters.

Blocks	Throughput (samples/s)	GFlops (Billion)	Parameters (M)
6-blocks	19.1	1.74	160.9
5-blocks	21.8 (+2.7)	1.45 (-0.29)	134.5 (-16%)
4-blocks	25.6 (+6.5)	1.16 (-0.58)	108.2 (-32%)
3-blocks	30.8 (+11.7)	0.87 (-0.87)	81.9 (-49%)
2-blocks	38.0 (+18.9)	0.59 (-1.15)	55.5 (-65%)

Inference Efficiency: To assess inference efficiency, we use the QMOF dataset, which has a large average graph size, making it a representative benchmark for evaluating computational cost. Using a V100 GPU, we evaluate inference throughput on a subset of the QMOF validation set and report results in Table 2.

Reducing the number of interaction blocks improves inference throughput by lowering computational overhead. The 6-block model processes 19.1 samples per second, while reducing to 5 blocks increases throughput to 21.8 samples per second, reflecting a moderate efficiency gain. Further reductions yield more substantial improvements: the 4-block model achieves 25.6 samples per second, providing a 1.3x speedup over the 6-block model while maintaining strong predictive performance. The 3-block and 2-block models offer even higher throughput, but at the cost of accuracy loss as seen in Figure 3.

Our block reduction strategy greatly lowers computational cost, as reflected in GFlops and parameter count. The 4-block model reduces GFlops from 1.74B to 1.16B and parameters from 160.9M to 108.2M (-32%), demonstrating its efficiency gains. This targeted reduction makes the block-reduced models particularly well-suited for deployment in resource-constrained environments and high-throughput screening applications, where maintaining predictive performance while optimizing computational efficiency is essential.

5. Generalization to Other Architectures

Our analysis in Section 4 is based on the GemNet-OC [18] architecture within the JMP [38] pre-training paradigm. To the best of our knowledge, JMP is the only large-scale, multi-domain pre-training approach for atomic property prediction. While several architectures provide pre-trained checkpoints on the OC20 catalysis dataset [8], it is unclear whether these models generalize as well as JMP across diverse downstream tasks.

To evaluate the generalizability of our block reduction methodology beyond GemNet-OC, we extend our study to EquiformerV2 [25], a state-of-the-art transformer-based equivariant model for atomic property prediction. Using the same JMP pre-training methodology, we pretrain EquiformerV2 with 31M parameters on the same four upstream tasks. Due to computational constraints, we limit pre-training to 12 A100 GPU days. We then apply our block reduction and knowledge distillation strategies and evaluate the model across multiple downstream tasks. While block reduction and knowledge distillation considerably reduced training time with minimal performance loss for JMP-L (160M parameters), their impact on EquiformerV2 (31M parameters) was more limited. This suggests that our approach is most effective for larger models, where over-parameterization leaves room for optimization. We include the evaluation results and discussion in the Appendix.

6. Conclusion

In this work, we explored strategies to enhance the efficiency of foundation models for atomic property prediction. By analyzing the role of individual layers in JMP-L, we found that deeper interaction blocks contribute less to predictive accuracy, making them suitable candidates for pruning. Our results show that reducing JMP-L’s parameter count by 32% improves inference throughput by 1.3x while maintaining comparable performance. Additionally, we demonstrated that knowledge distillation can help mitigate performance degradation in certain tasks. Our methodology reduces inference time and computational cost, making the optimized models practical for scientists conducting research in resource-constrained settings and high-throughput screening, where both accuracy and efficiency are crucial. We hope this study encourages further research into efficient training and inference for molecular property prediction, paving the way for lighter models in molecular and materials discovery.

Acknowledgments. This work is supported by the KAUST Center of Excellence for Generative AI under award number 5940. The computational resources are provided by IBEX, which is managed by the Supercomputing Core Laboratory at KAUST. Yasir is supported by Saudi Aramco.

References

- [1] Ilyes Batatia, Simon Batzner, Dávid Péter Kovács, Albert Musaelian, Gregor N. C. Simm, Ralf Drautz, Christoph Ortner, Boris Kozinsky, and Gábor Csányi. The design space of e(3)-equivariant atom-centered interatomic potentials, 2022. [1](#)
- [2] Ilyes Batatia, David Peter Kovacs, Gregor N. C. Simm, Christoph Ortner, and Gabor Csanyi. MACE: Higher order equivariant message passing neural networks for fast and accurate force fields. In *Advances in Neural Information Processing Systems*, 2022. [1](#)
- [3] Ilyes Batatia, Philipp Benner, Yuan Chiang, Alin M Elena, Dávid P Kovács, Janosh Riebesell, Xavier R Advincula, Mark Asta, William J Baldwin, Noam Bernstein, et al. A foundation model for atomistic materials chemistry. *arXiv preprint arXiv:2401.00096*, 2023. [1](#), [2](#)
- [4] Jörg Behler and Michele Parrinello. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Physical Review Letters*, 98(14), 2007. [1](#)
- [5] Davis Blalock, Jose Javier Gonzalez Ortiz, Jonathan Frankle, and John Guttag. What is the state of neural network pruning? *Proceedings of machine learning and systems*, 2: 129–146, 2020. [2](#)
- [6] Mihail Bogojeski, Leslie Vogt-Maranto, Mark E. Tuckerman, Klaus-Robert Müller, and Kieron Burke. Quantum chemical accuracy from density functional approximations via machine learning. *Nature Communications*, 11(1), 2020. [1](#)
- [7] Cristian Bucilua, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, page 535–541, New York, NY, USA, 2006. Association for Computing Machinery. [2](#)
- [8] Lowik Chanussot, Abhishek Das, Siddharth Goyal, Thibaut Lavril, Muhammed Shuaibi, Morgane Riviere, Kevin Tran, Javier Heras-Domingo, Caleb Ho, Weihua Hu, et al. Open catalyst 2020 (oc20) dataset and community challenges. *Acs Catalysis*, 11(10):6059–6072, 2021. [1](#), [5](#), [8](#)
- [9] Hongrong Cheng, Miao Zhang, and Javen Qinfeng Shi. A survey on deep neural network pruning-taxonomy, comparison, analysis, and recommendations, 2024. [2](#)
- [10] Stefan Chmiela, Alexandre Tkatchenko, Huziel E Sauceda, Igor Poltavsky, Kristof T Schütt, and Klaus-Robert Müller. Machine learning of accurate energy-conserving molecular force fields. *Science advances*, 3(5):e1603015, 2017. [6](#)
- [11] Stefan Chmiela, Valentin Vassilev-Galindo, Oliver T Unke, Adil Kabylda, Huziel E Sauceda, Alexandre Tkatchenko, and Klaus-Robert Müller. Accurate global machine learning force fields for molecules with hundreds of atoms. *Science Advances*, 9(2):eadf0873, 2023. [6](#)
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. [2](#)
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. [2](#)
- [14] Alexander Dunn, Qi Wang, Alex Ganose, Daniel Dopp, and Anubhav Jain. Benchmarking materials property prediction methods: the matbench test set and automatminer reference algorithm. *npj Computational Materials*, 6(1):138, 2020. [1](#)
- [15] Peter Eastman, Pavan Kumar Behara, David L Dotson, Raimondas Galvelis, John E Herr, Josh T Horton, Yuezhi Mao, John D Chodera, Benjamin P Pritchard, Yuanqing Wang, et al. Spice, a dataset of drug-like molecules and peptides for training machine learning potentials. *Scientific Data*, 10(1):11, 2023. [1](#), [6](#)
- [16] Filip Ekström Kelvinius, Dimitar Georgiev, Artur Toshev, and Johannes Gasteiger. Accelerating molecular graph neural networks via knowledge distillation. *Advances in Neural Information Processing Systems*, 36, 2024. [2](#), [3](#), [5](#)
- [17] Johannes Gasteiger, Florian Becker, and Stephan Günnemann. Gemnet: Universal directional graph neural networks for molecules. *Advances in Neural Information Processing Systems*, 34:6790–6802, 2021. [3](#)
- [18] Johannes Gasteiger, Muhammed Shuaibi, Anuroop Sriram, Stephan Günnemann, Zachary Ulissi, C Lawrence Zitnick, and Abhishek Das. Gemnet-oc: developing graph neural networks for large and diverse molecular simulation datasets. *arXiv preprint arXiv:2204.02782*, 2022. [2](#), [3](#), [8](#)
- [19] Andrey Gromov, Kushal Tirumala, Hassan Shapourian, Paolo Glorioso, and Daniel A. Roberts. The unreasonable ineffectiveness of the deeper layers, 2024. [2](#), [3](#)
- [20] Gousia Habib, Tausifa Jan Saleem, and Brejesh Lall. Knowledge distillation in vision transformers: A critical review, 2024. [2](#)
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [2](#)
- [22] Geoffrey Hinton. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. [2](#), [3](#)
- [23] Adeesh Kolluru, Muhammed Shuaibi, Aini Palizhati, Nima Shoghi, Abhishek Das, Brandon Wood, C. Lawrence Zitnick, John R Kitchin, and Zachary W Ulissi. Open challenges in developing generalizable large scale machine learning models for catalyst discovery, 2022. [1](#)
- [24] Dávid Péter Kovács, J Harry Moore, Nicholas J Browning, Ilyes Batatia, Joshua T Horton, Venkat Kapil, Ioan-Bogdan Magdău, Daniel J Cole, and Gábor Csányi. Maceoff23: Transferable machine learning force fields for organic molecules. *arXiv preprint arXiv:2312.15211*, 2023. [2](#)
- [25] Yi-Lun Liao, Brandon Wood, Abhishek Das, and Tess Smidt. Equiformerv2: Improved equivariant transformer for scaling to higher-degree representations. *arXiv preprint arXiv:2306.12059*, 2023. [8](#), [11](#)
- [26] Chuang Liu, Xueqi Ma, Yibing Zhan, Liang Ding, Dapeng Tao, Bo Du, Wenbin Hu, and Danilo Mandic. Comprehensive graph gradual pruning for sparse training in graph neural networks, 2022. [2](#)

- [27] Xinyin Ma, Gongfan Fang, and Xinchao Wang. LLM-pruner: On the structural pruning of large language models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. 3
- [28] Jens K. Nørskov, Frank Abild-Pedersen, Felix Studt, and Thomas Bligaard. Density functional theory in surface chemistry and catalysis. *Proceedings of the National Academy of Sciences*, 108(3):937–943, 2011. 1
- [29] Raghunathan Ramakrishnan, Pavlo O Dral, Matthias Rupp, and O Anatole Von Lilienfeld. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific data*, 1(1): 1–7, 2014. 6
- [30] Andrew S Rosen, Shaelyn M Iyer, Debmalya Ray, Zhen-peng Yao, Alan Aspuru-Guzik, Laura Gagliardi, Justin M Notestein, and Randall Q Snurr. Machine learning the quantum-chemical properties of metal–organic frameworks for accelerated materials discovery. *Matter*, 4(5):1578–1597, 2021. 6
- [31] Andrew S. Rosen, Victor Fung, Patrick Huck, Cody T. O’Donnell, Matthew K. Horton, Donald G. Truhlar, Kristin A. Persson, Justin M. Notestein, and Randall Q. Snurr. High-throughput predictions of metal–organic framework electronic properties: theoretical challenges, graph neural networks, and data exploration. *npj Computational Materials*, 8(1), 2022. 1
- [32] Victor T. Sabe, Thandokuhle Ntombela, Lindiwe A. Jhamba, Glenn E.M. Maguire, Thavendran Govender, Tricia Naicker, and Hendrik G. Kruger. Current trends in computer aided drug design and a highlight of drugs discovered via computational techniques: A review. *European Journal of Medicinal Chemistry*, 224:113705, 2021. 1
- [33] Mathias Schreiner, Arghya Bhowmik, Tejs Vegge, Jonas Busk, and Ole Winther. Transition1x - a dataset for building generalizable reactive machine learning potentials. *Scientific Data*, 9(1), 2022. 1
- [34] Mathias Schreiner, Arghya Bhowmik, Tejs Vegge, Jonas Busk, and Ole Winther. Transition1x-a dataset for building generalizable reactive machine learning potentials. *Scientific Data*, 9(1):779, 2022. 5
- [35] Kristof Schütt, Pieter-Jan Kindermans, Huziel Enoc Saucedo Felix, Stefan Chmiela, Alexandre Tkatchenko, and Klaus-Robert Müller. Schnet: A continuous-filter convolutional neural network for modeling quantum interactions. *Advances in neural information processing systems*, 30, 2017. 3
- [36] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 618–626, 2017. 3
- [37] Vinit Sharma, Chenchen Wang, Robert G. Lorenzini, Rui Ma, Qiang Zhu, Daniel W. Sinkovits, Ghanshyam Pili-ania, Artem R. Oganov, Sanat Kumar, Gregory A. Sotzing, Steven A. Boggs, and Rampi Ramprasad. Rational design of all organic polymer dielectrics. *Nature Communications*, 5(1), 2014. 1
- [38] Nima Shoghi, Adeesh Kolluru, John R Kitchin, Zachary W Ulissi, C Lawrence Zitnick, and Brandon M Wood. From molecules to materials: Pre-training large generalizable models for atomic property prediction. *arXiv preprint arXiv:2310.16802*, 2023. 1, 2, 5, 6, 8, 11
- [39] Sietsma and Dow. Neural net pruning-why and how. In *IEEE 1988 International Conference on Neural Networks*, pages 325–333 vol.1, 1988. 2
- [40] Justin S Smith, Roman Zubatyuk, Benjamin Nebgen, Nicholas Lubbers, Kipton Barros, Adrian E Roitberg, Olexandr Isayev, and Sergei Tretiak. The ani-1ccx and ani-1x data sets, coupled-cluster and density functional theory properties for molecules. *Scientific data*, 7(1):134, 2020. 1, 5
- [41] Anuroop Sriram, Sihoon Choi, Xiaohan Yu, Logan M. Brabson, Abhishek Das, Zachary Ulissi, Matt Uyttendaele, Andrew J. Medford, and David S. Sholl. The open dac 2023 dataset and challenges for sorbent discovery in direct air capture, 2023. 1
- [42] Mingjie Sun, Zhuang Liu, Anna Bair, and J. Zico Kolter. A simple and effective pruning approach for large language models, 2024. 2
- [43] Yijun Tian, Shichao Pei, Xiangliang Zhang, Chuxu Zhang, and Nitesh V. Chawla. Knowledge distillation on graphs: A survey, 2023. 2
- [44] Richard Tran, Janice Lan, Muhammed Shuaibi, Brandon M Wood, Siddharth Goyal, Abhishek Das, Javier Heras-Domingo, Adeesh Kolluru, Ammar Rizvi, Nima Shoghi, et al. The open catalyst 2022 (oc22) dataset and challenges for oxide electrocatalysts. *ACS Catalysis*, 13(5):3066–3084, 2023. 1, 5
- [45] Xiaohan Xu, Ming Li, Chongyang Tao, Tao Shen, Reynold Cheng, Jinyang Li, Can Xu, Dacheng Tao, and Tianyi Zhou. A survey on knowledge distillation of large language models, 2024. 2
- [46] Fang Yu, Kun Huang, Meng Wang, Yuan Cheng, Wei Chu, and Li Cui. Width & depth pruning for vision transformers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3143–3151, 2022. 2
- [47] Sheheryar Zaidi, Michael Schaarschmidt, James Martens, Hyunjik Kim, Yee Whye Teh, Alvaro Sanchez-Gonzalez, Peter Battaglia, Razvan Pascanu, and Jonathan Godwin. Pre-training via denoising for molecular property prediction, 2022. 1
- [48] Liang Zeng, Lanqing Li, and Jian Li. Molkd: Distilling cross-modal knowledge in chemical reactions for molecular property prediction, 2023. 2
- [49] Mingyang Zhang, Hao Chen, Chunhua Shen, Zhen Yang, Linlin Ou, Xinyi Yu, and Bohan Zhuang. Loraprune: Structured pruning meets low-rank parameter-efficient fine-tuning, 2024. 2
- [50] Gengmo Zhou, Zhifeng Gao, Qiankun Ding, Hang Zheng, Hongteng Xu, Zhewei Wei, Linfeng Zhang, and Guolin Ke. Uni-mol: A universal 3d molecular representation learning framework. In *The Eleventh International Conference on Learning Representations*, 2023. 1

Appendix

A. Beyond GemNet-OC

In this section, we extend our block reduction and knowledge distillation methodology to EquiformerV2 [25], a state-of-the-art transformer-based equivariant model for atomic property prediction. Following the same pipeline as in the main paper, we first apply GradCAM as described in Section 3.2 to assess interaction block importance. We then pre-train EquiformerV2 following the same methodology used for JMP before evaluating its performance on the downstream tasks.

A.1. Block Relevance

Using the JMP [38] pre-training datasets (OC20, OC22, ANI-1x, and Transition-1x), we perform forward passes on 1,000 samples from each dataset to extract features from the EquiformerV2 backbone (31M parameters). This backbone consists of 8 transformer blocks. We save the output features of each block and apply GradCAM (as described in Section 3.2) to assess the relevance of each block to the final output prediction.

Appendix Figure 1 presents the relevance scores of EquiformerV2’s transformer blocks. Similar to GemNet-OC, we observe a diminishing return in later blocks, where deeper layers contribute less to the final prediction. This trend suggests that the later transformer blocks are strong candidates for pruning, reinforcing the effectiveness of our block reduction strategy for improving model efficiency.

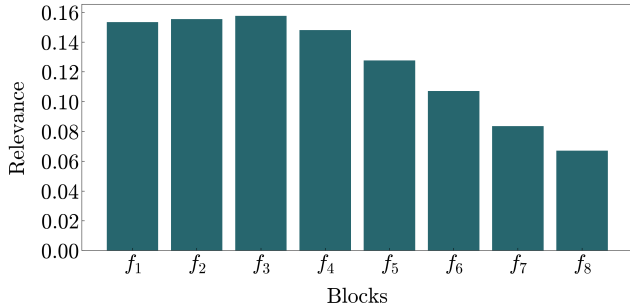


Figure 1. **Block Relevance Analysis for EquiformerV2.** This figure illustrates the contribution of each transformer block in EquiformerV2 to the final prediction. The model consists of eight transformer blocks, f_1 to f_8 , with relevance decreasing in the deeper layers. The results suggest that later blocks are strong candidates for pruning to improve efficiency with minimal performance loss.

A.2. Block Reduction and Distillation in Pre-training for EquiformerV2

Pre-training setting: While GemNet-OC has publicly available pre-trained weights on all four JMP datasets,

EquiformerV2 only provides pre-trained weights for OC20 or OC22 individually. To ensure a fair and consistent comparison, we initialize the EquiformerV2 backbone with the FAIRChem pre-trained weights (EquiformerV2-31M-S2EF-OC20-All+MD). We then pre-train the model following the JMP pre-training paradigm, using a uniform balancing strategy with 2M samples from each upstream dataset, totaling 8M samples. The pre-training is conducted over 2 epochs using 4 A100 GPUs, with a total compute time of 12 A100 GPU-days.

After pre-training, we apply block reduction (BR) and knowledge distillation (KD) following the approach outlined in Section 4.2 of the main paper. For block reduction, we sequentially remove later blocks, constructing smaller models with 7, 6, 5, and 4 transformer blocks. Since removing blocks may lead to performance degradation, we conduct a brief KD phase to recover performance for each pruned model. To manage computational cost, we limit the KD phase to 1.5% of the pre-training size. This process requires approximately 63 A100 GPU-hours.

Observations: Appendix Table 1 presents the force MAE performance across various pre-training tasks for EquiformerV2. Consistent with GemNet-OC results in the main paper, we observe that removing later blocks leads to a proportional performance drop, confirming that deeper layers contribute less to the final prediction. Applying knowledge distillation after block reduction considerably improves performance, mitigating much of the accuracy loss. For example, in the 4-block model, KD reduces the performance gap with the teacher model from -50.9 meV/Å to -5.9 meV/Å for OC20 and from -40.4 meV/Å to -10.3 meV/Å for OC22. Similar gains are observed for ANI-1x and Transition-1x, highlighting the effectiveness of KD in compensating for the loss of capacity from pruning.

These results reinforce the effectiveness of our block reduction and KD approach across different model architectures, demonstrating that smaller, more efficient models can retain strong predictive accuracy with minimal fine-tuning overhead. Next, we evaluate the performance of these pruned models on downstream tasks.

A.3. Evaluation on Downstream Tasks using EquiformerV2

We evaluate the impact of our block reduction and pre-training knowledge distillation on downstream tasks, following the fine-tuning setup in Section 4.3. Distillation was applied only during pre-training to retain useful representations before fine-tuning.

Appendix Figure 2 presents the results across multiple tasks. In Spice and QM9, the 7-block models achieve comparable or better performance than the 8-block teacher model, demonstrating that EquiformerV2 can be effectively compressed while maintaining accuracy for some tasks. In

Table 1. **Force MAE Evaluation During Pre-training of EquiformerV2.** We analyze the impact of block reduction (BR) on pre-training performance for EquiformerV2, comparing results with and without knowledge distillation (KD). The table reports force MAE (meV/Å) across the OC20, OC22, ANI-1x, and Transition-1x datasets. Similar to GemNet-OC, we observe that later blocks contribute less to the final prediction, and KD mitigates performance degradation relative to the full model.

	# of Blocks	OC20	OC22	ANI-1x	Transition-1x
BR	4	77.0 (-50.9)	71.5 (-40.4)	185.5 (-142.4)	69.6 (-52.0)
+ KD	4	32.0 (-5.9)	41.4 (-10.3)	68.8 (-25.7)	26.9 (-9.3)
BR	5	51.1 (-25.0)	54.1 (-23.0)	85.4 (-42.3)	40.8 (-23.2)
+ KD	5	30.3 (-4.2)	38.2 (-7.1)	64.2 (-21.1)	24.3 (-6.7)
BR	6	35.2 (-9.1)	40.9 (-9.8)	60.1 (-17.0)	25.6 (-8.0)
+ KD	6	29.2 (-3.1)	35.3 (-4.2)	58.0 (-14.9)	21.8 (-4.2)
BR	7	27.9 (-1.8)	33.4 (-2.3)	42.0 (1.1)	19.9 (-2.3)
+ KD	7	27.4 (-1.3)	32.8 (-1.7)	53.4 (-10.3)	19.7 (-2.1)
EquiformerV2 (Teacher)	8	26.1	31.1	43.1	17.6

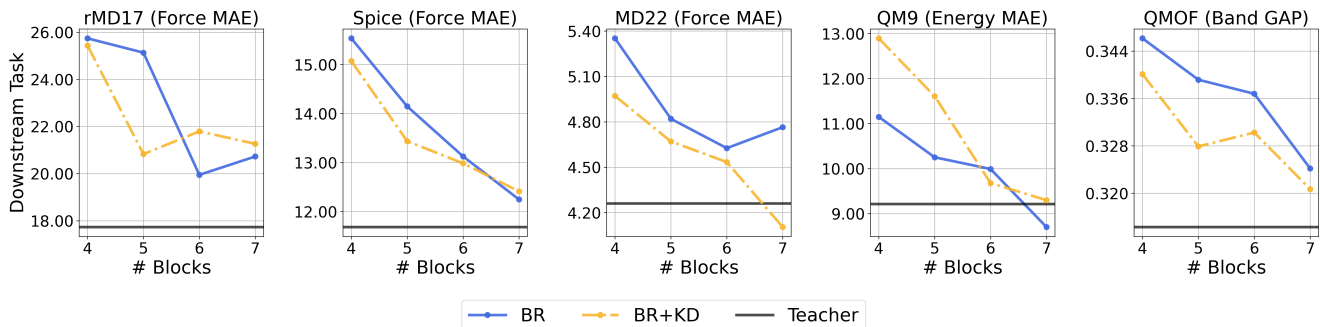


Figure 2. **Evaluation on downstream tasks using EquiformerV2.** We assess the performance of EquiformerV2 across various downstream tasks after applying block reduction (BR) and knowledge distillation (KD). Performance is reported in meV/Å for force targets, meV for the QM9 energy target, and eV for the QMOF band gap target. The teacher model refers to the original EquiformerV2 with 8 transformer blocks.

contrast, rMD17 and QMOF show a decline in performance as blocks are removed, even when distillation is applied. This suggests that the 31M-parameter teacher model was already compact, leaving little room for pruning without sacrificing accuracy.

Distillation is most beneficial at 4 and 5 blocks, considerably recovering performance lost due to pruning. However, for 6 and 7 blocks, the gains from distillation diminish, as block reduction alone is sufficient for maintaining accuracy—except in MD22, where distillation leads to improved results over the full model.

These results highlight that while block reduction and knowledge distillation were effective for GemNet-OC, their impact on EquiformerV2 is mixed. One possible reason is that our pre-training duration for EquiformerV2 was limited to 8M samples compared to 120M samples in JMP, con-

strained by computational resources, which may have prevented the models from fully capturing useful representations. Additionally, EquiformerV2-31M might be too small to benefit from block reduction, as reducing blocks further constrains its capacity.

B. Training from Scratch

In the main paper, all experiments use models initialized from the JMP-L pre-trained checkpoint. Specifically, we apply block reduction to create smaller backbones by pruning layers from the full model—for example, constructing a 5-block GemNet-OC by removing the final block from a 6-block pre-trained backbone. But what if we instead train the 5-block model from scratch? How does this compare to pruning a pre-trained full-sized model?

Appendix Figure 3 compares the downstream perfor-

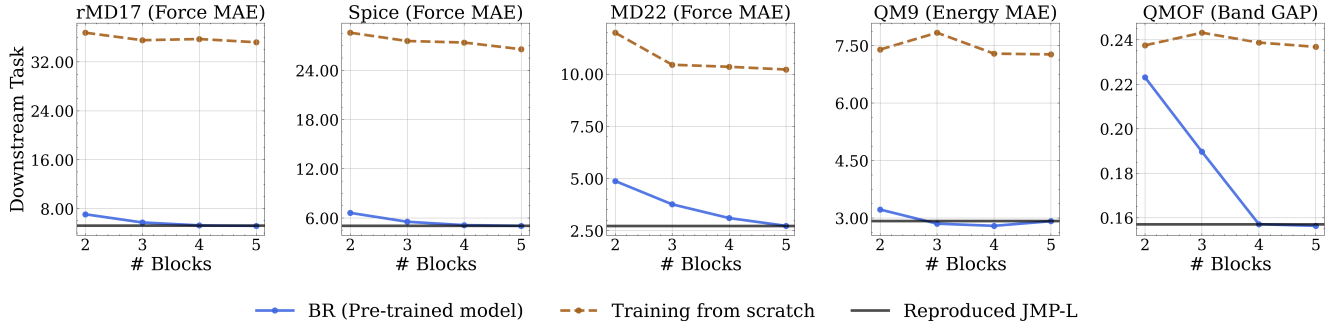


Figure 3. **Training from scratch vs. block reduction of a pre-trained model.** Validation performance of smaller models across five downstream tasks when initialized from scratch (brown dashed line) versus when derived through block reduction from a pre-trained 6-block JMP-L model (blue line). The results highlight the importance of pre-training, as reduced models retain significantly better performance when initialized from a larger pre-trained backbone.

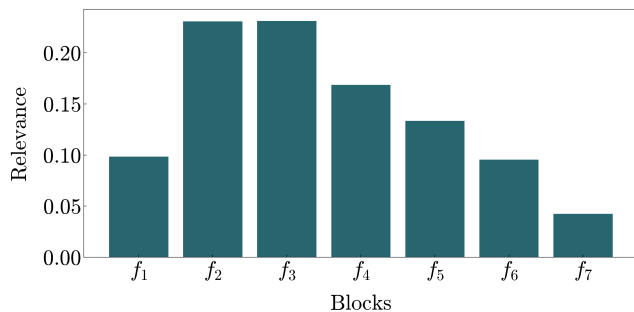


Figure 4. **Block Relevance Analysis (Shown in Main Paper).** GradCAM-based importance scores for each block in JMP-L, previously presented in the main paper. The first block, f_1 , represents the embedding output, while f_2 to f_7 correspond to the six interaction blocks of GemNet-OC.

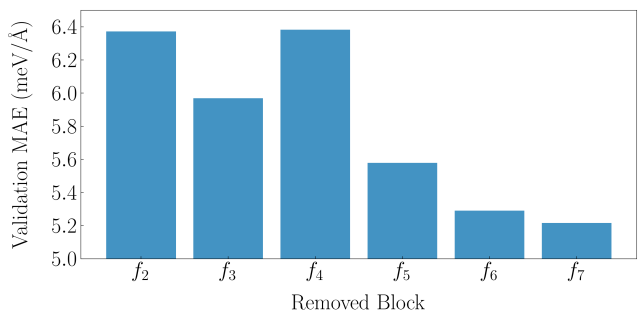


Figure 5. **Layer-Wise Ablation Study.** Validation MAE on rMD17 when individually removing each interaction block (f_2 to f_7) from the 6-block pre-trained GemNet-OC model. The results show that removing later blocks leads to smaller performance loss, supporting their lower importance as suggested by GradCAM.

mance of models trained from scratch (brown dashed line) with those obtained via block reduction from a pre-trained model (blue line). Across all tasks, training from scratch consistently underperforms the block-reduced counterparts, often by a large margin. These results highlight the effectiveness of our block reduction strategy and demonstrate the importance of leveraging pre-trained weights, even when deploying smaller backbones.

C. Layer-by-Layer Analysis of Block Reduction

As discussed in Section 3.2 of the main paper, our block reduction strategy is guided by block importance scores obtained using GradCAM. For instance, to create a 5-block GemNet-OC backbone, we remove the least important block—identified as the final block (f_7). However, this approach assumes that GradCAM accurately reflects functional relevance. Yet, removing a different block might lead to comparable downstream performance, which would cast doubt on the reliability of GradCAM for guiding pruning decisions.

To evaluate the validity of the GradCAM scores, we conduct a layer-by-layer ablation, removing one interaction block at a time—starting from the first non-embedding layer (f_2) through to the final interaction block (f_7)—resulting in six separate runs. We note that removing f_7 matches the 5-block configuration used in the main paper. To manage computational cost, we restrict this analysis to the rMD17 dataset and fix the backbone size to 5 blocks.

In Appendix Figures 4 and 5 we present the block relevance scores and the corresponding downstream performance from the layer-by-layer ablation analysis, respectively. As predicted by GradCAM, removing later blocks leads to less performance degradation compared to earlier ones. Overall, the trend in performance drop from the first interaction block (f_2) to the last (f_7) aligns well with the relevance scores—except for f_3 and f_4 , where removing f_4 leads to a larger drop than f_3 . This suggests that GradCAM provides more reliable importance estimates for deeper blocks closer to the prediction head.