Spherical Channels for Modeling Atomic Interactions

Anonymous Author(s)
Affiliation
Address
email

Abstract

Modeling the energy and forces of atomic systems is a fundamental problem in computational chemistry with the potential to help address many of the world's most pressing problems, including those related to energy scarcity and climate change. These calculations are traditionally performed using Density Functional Theory, which is computationally very expensive. Machine learning has the potential to dramatically improve the efficiency of these calculations from days or hours to seconds.

We propose the Spherical Channel Network (SCN) to model atomic energies and forces. The SCN is a graph neural network where nodes represent atoms and edges their neighboring atoms. The atom embeddings are a set of spherical functions, called spherical channels, represented using spherical harmonics. We demonstrate, that by rotating the embeddings based on the 3D edge orientation, more information may be utilized while maintaining the rotational equivariance of the messages. While equivariance is a desirable property, we find that by relaxing this constraint in both message passing and aggregation, improved accuracy may be achieved. We demonstrate state-of-the-art results on the large-scale Open Catalyst dataset in both energy and force prediction for numerous tasks and metrics.

1 Introduction

Modeling the properties of atomic systems is a foundational challenge in computational chemistry and critical to advancing technologies across numerous application domains. Notable applications include drug discovery [28, 35] and the design of new catalysts for renewable energy storage to help in addressing climate change [44, 29]. For catalyst discovery, new materials are currently evaluated using Density Functional Theory (DFT) that can estimate atomic energies and forces but is computationally very expensive; taking hours or days to evaluate a single material. Machine Learning (ML) has the potential to approximate DFT and dramatically speed up these calculations, allowing for high throughput screening of new materials to help address some of the world's most pressing challenges.

Our goal is to approximate DFT calculations using ML. An ML model takes as input a set of atom positions and their atomic numbers. As outputs, a model calculates the structure's energy (or other properties) and the per-atom forces, i.e., the forces exerted on each atom by the other atoms. A common approach to this problem is to use Graph Neural Networks (GNNs) [15, 43] where each node represents an atom and the set of nearby atoms as edges [33, 13, 19, 31, 34, 41, 27, 21].

A key challenge in network design is balancing the use of model constraints. When modeling atomic systems, a common constraint is SO(3) rotation equivariance [39, 2, 1, 38, 30, 32], *i.e.*, if the atomic

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

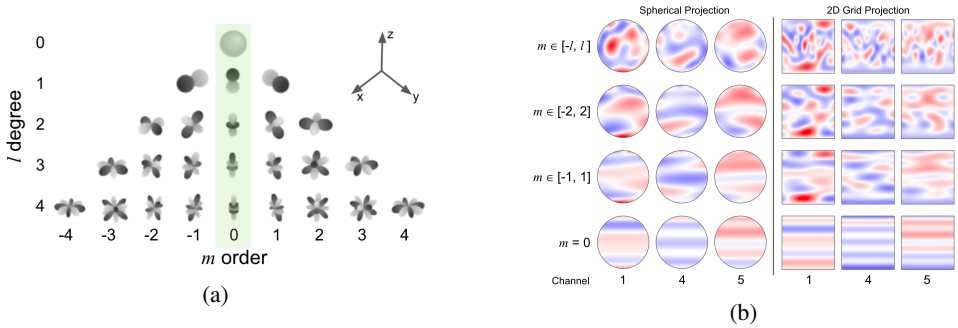

Figure 1: (a) Illustration of spherical harmonics up to $l = 4$. Note the $m = 0$ bases are symmetric about the z-axis (highlighted in green). (b) Visualization of 3 spherical channels with $l = 8$ when varying the number of orders: all, $m \in [-2, 2]$, $m \in [-1, 1]$, $m = 0$. Note how the resolution in the x and y dimensions decreases with fewer m until values are constant for a given z for $m = 0$. Spherical projections are shown on the left (only half of the channel is visible) and a 2D grid projection (polar and azimuthal) of the same channels is shown on the right.

system is rotated, the energies should remain constant and the atomic forces should similarly rotate. While this provides strong priors on the model to help in generalization, especially for smaller datasets[6, 28], it can result in limiting the expressiveness of the network due to restrictions on non-linear transformations [2, 38, 30, 32], or limiting interactions to pairs [31, 34, 41], triplets [21] or quadruplets of atoms [25, 22, 11]. Alternatively, a non-equivariant model [18] can provide more freedom to the model, but lead to the model needing to learn approximate equivariance through methods such as data augmentation [36]. To draw an analogy with image detection, the use of a CNN [24] provides translation equivariance and removes the need for the network to learn how to detect the same object at different locations. However, most CNNs are not equivariant to scale or rotation [40, 39], but are still effective in learning approximate equivariance through rotation and scale diversity in the training data. What is the analogous balance of constraints for modeling atomic systems?

In this paper, we propose a GNN [15, 43] that balances the use of model constraints to aid in generalization while providing the network with the flexibility to learn accurate representations. We introduce the Spherical Channel Network (SCN) that explicitly models relative orientations of all neighboring atoms; information that is critical to accurately predicting atomic properties. Each node's embedding is a set of functions defined on the surface of a sphere ($S^2 \rightarrow \mathbb{R}$). The functions are represented using spherical harmonics, similar to approaches that build strictly equivariant models [38, 1, 2, 9]. During message passing, angular information between atoms is conveyed by rotating or steering [4] the embeddings based on each edge's orientation [37, 25]. We identify an expanded set of spherical harmonic coefficients that are invariant to rotation, which can provide rich information while maintaining a message's rotation equivariance. In addition, we demonstrate that if the equivariance constraint is relaxed, improved performance can be achieved by using additional coefficients. We further improve the expressivity of the network by performing a pointwise non-linear function on the embeddings during message aggregation.

We demonstrate our Spherical Channel Network on the large-scale Open Catalyst 2020 (OC20) dataset [5], which contains atomic structures useful for numerous applications important to addressing climate change. State-of-the-art results are achieved for atomic force and initial structure to relaxed energy prediction with improvements of $8\% - 11\%$. In addition, we demonstrate our model is more sample efficient as compared to other state-of-the-art models.

2 Approach

Given an atomic structure with n atoms, our goal is to predict the structure's energy E and the per-atom forces $\boldsymbol{f}_i$ for each atom $i \in n$. These values are estimated using a Graph Neural Network

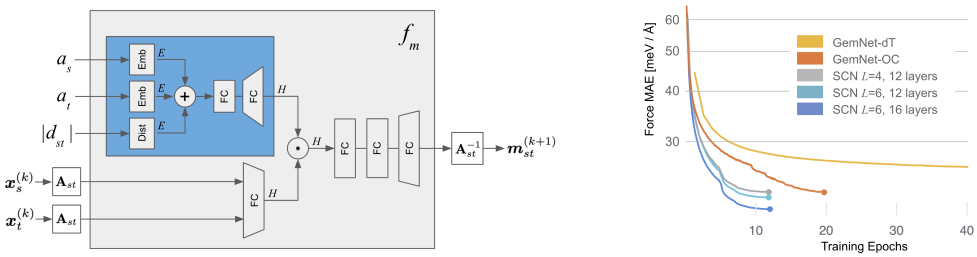

(a) (b)

Figure 2: (a) Block diagram of message passing function $\boldsymbol{F}_m$ in Equation 2 for source atom s to target atom t. The atomic numbers a_s and a_t, distance between the atoms $|d_{st}|$, and embeddings $\boldsymbol{x}_s^{(k)}$ and $\boldsymbol{x}_t^{(k)}$ are given as input. (b) Training curves for SCN and GemNet models for force MAEs per epoch evaluated on val ID 30k. Note how the SCN model is significantly more sample efficient during training.

(GNN) [15, 43] where each node represents an atom and edges represent nearby atoms. As input, the network is given the distance d_{ij} between atoms i and j, and each atom's atomic number a_i. The neighbors N_i for an atom i are determined using a fixed distance threshold, or by picking a fixed number of closest atoms.

2.1 Node Embeddings

The angular or relative orientation information between atoms is critical to accurately modeling atomic properties such as energies and forces [21, 22]. Inspired by this, our node embedding explicitly models the angular information from all neighboring atoms. Each node i's embedding s_i is a set of C functions or channels represented on a sphere ($S^2 \to \mathbb{R}$), whose argument is a 3D unit vector indicating the orientation. That is, $s_{ic}(\hat{\boldsymbol{r}})$ is the value of channel c for node i for some orientation $\hat{\boldsymbol{r}} \in \mathbb{R}^3$. Since the spherical channels contain orientation information over the entire sphere, the network may reason about directional information for all neighboring atoms and not just atom pairs, triplets, etc. Each spherical channel c may be represented using several different approaches, such as discrete 2D grids sampled over the sphere, or spherical harmonics. As we describe later, we use spherical harmonics due to their property of being SO(3)-equivariant (3D rotation equivariant).

Using spherical harmonics, a function $s_c(\hat{\boldsymbol{r}})$ defined on the sphere is represented by a set of weighted spherical harmonics Y_l^m with $s_c(\hat{\boldsymbol{r}}) = \sum_{l,m} \boldsymbol{x}_{lmc} Y_l^m(\hat{\boldsymbol{r}})$ where l and m are the degree and order of the basis functions with $m \in [-l, l]$ and $l \in L$. $\boldsymbol{x}_{lmc}$ are the spherical harmonic coefficients, which we refer to as the coefficients of the spherical channel. For every degree l there exists $2l + 1$ spherical harmonics (see Figure 1(a) for an illustration of functions for $l \leq 4$), which results in a total of $(L + 1)^2$ basis functions and coefficients up to degree L. Therefore for a maximum degree of L with C channels, $\boldsymbol{x}$ has size $(L + 1)^2 \times C$.

The spherical channels are updated by the GNN through message passing for K layers to obtain the final node embeddings $S^{(K)}$. $S^{(0)}$ is initialized from an embedding based on the atom's atomic number a_i for $l = 0$ coefficients and the $l \neq 0$ coefficients are set to zero. The nodes' embeddings are updated by first calculating a set of messages m_{ij} for each edge, which are then aggregated at each node. Finally, the energy and forces are estimated from $S^{(K)}$. We describe each of these steps in turn.

2.2 Message Passing

Given a target node t and its neighbors $s \in N_t$ we want to update the embeddings $\boldsymbol{x}_t^{(k)}$ at iteration $k \in K$. The embeddings $\boldsymbol{x}_t^{(k)}$ are a set of spherical harmonic coefficients indexed by their degree l, order m and channel c. An important and useful property of spherical harmonics is the function represented by the coefficients is steerable [10, 4], *i.e.*, it can be rotated in 3D space using a linear

transformation of the coefficients. Specifically, for a 3D rotation matrix $\mathbf{R}$ there exists a matrix called a Wigner D-matrix A_l of size $(2l+1 \times 2l+1)$ that rotates the coefficients of degree l by the rotation $\mathbf{R}$. If $\boldsymbol{x}_{ilc}$ are all coefficients of degree l for node i and channel c, then for any orientation $\hat{\boldsymbol{r}}$:

$$s_{ic}(\mathbf{R}\hat{\boldsymbol{r}}) = \sum_{l}^{L} \boldsymbol{x}_{ilc} \cdot Y_l(\mathbf{R}\hat{\boldsymbol{r}}) = \sum_{l}^{L} (\mathbf{A}_l \boldsymbol{x}_{ilc}) \cdot Y_l(\hat{\boldsymbol{r}}). \tag{1}$$

When calculating the message m_{st} from atom s to atom t, we want to use the information contained in both $\boldsymbol{x}_s^{(k)}$ and $\boldsymbol{x}_t^{(k)}$ given the context of the edge's orientation $\hat{\boldsymbol{d}}_{st}$, $\hat{\boldsymbol{d}}_{st} = \boldsymbol{d}_{st}/|\boldsymbol{d}_{st}|$. We do this by rotating the embeddings by $\mathbf{R}_{st}$ for which $\mathbf{R}_{st}\hat{\boldsymbol{d}}_{st} = [0, 0, 1]^\top$, *i.e.*, the direction of the edge's unit vector $\hat{\boldsymbol{d}}_{st}$ is aligned with the z-axis. Thus when the messages are calculated, the problem is simplified since the atoms are always oriented along the z-axis. After calculating the messages using a neural network $\boldsymbol{F}_m$, they are rotated back to the global coordinate frame:

$$\boldsymbol{m}_{st}^{(k+1)} = \mathbf{A}_{st}^{-1} \boldsymbol{F}_m(\mathbf{A}_{st}\boldsymbol{x}_s^{(k)}, \mathbf{A}_{st}\boldsymbol{x}_t^{(k)}, |\boldsymbol{d}_{st}|, a_s, a_t), \tag{2}$$

where the matrices $\mathbf{A}_{st}$ and $\mathbf{A}_{st}^{-1}$ perform the rotation on the embeddings' coefficients corresponding to the 3D rotation matrices $\mathbf{R}_{st}$ and $\mathbf{R}_{st}^{-1}$ respectively. To simplify notation, we assume the coefficients in $\boldsymbol{x}_s$ across degree l and order m are flattened, and $\mathbf{A}_{st} \in \mathbb{R}^{(L+1)^2} \times \mathbb{R}^{(L+1)^2}$ is a block diagonal matrix constructed from the set of Wigner D-matrices across degrees. $\mathbf{A}_{st}^{-1}$ is the inverse of $\mathbf{A}_{st}$, with $\mathbf{A}_{st}^{-1} = \mathbf{A}_{st}^\top$. In addition to the rotated embeddings, $\boldsymbol{F}_m$ is provided as input edge information that is invariant to rotations; $|\boldsymbol{d}_{st}|$ is the magnitude of the distance between atoms s and t, and a_s and a_t their respective atomic numbers. Note that the rotation matrix $\mathbf{R}_{st}$ is not unique, since the roll rotation around the vector $\boldsymbol{d}_{st}$ is not specified and randomly sampled during training. The implications of this are discussed below.

The message function $\boldsymbol{F}_m$ is computed using a neural network as illustrated in Figure 2(a). The atomic numbers a_s and a_t are used to look up two independent embeddings of size $E = 128$, and a set of 1D basis functions are used to represent $|\boldsymbol{d}_{st}|$ using equally spaced Gaussians every 0.02 Å from 0 to the 8 Å with $\sigma = 0.04$ followed by a linear layer with E outputs. Their values are added together and passed through a neural network to produce a vector of size H. The rotated embeddings $\mathbf{A}_{st}\boldsymbol{x}_s^{(k)}$ and $\mathbf{A}_{st}\boldsymbol{x}_t^{(k)}$ are concatenated and passed through a single layer neural network to produce another vector of size H. These two are multiplied together to combine the edge information with the information contained in the rotated embeddings. Two more fully connected layers are performed each followed by a SiLU non-linear activation function [17]. Finally, a single linear layer is used to expand the output to the original size of the embeddings.

Note as stated before, the matrix $\mathbf{A}_{st}$ is not unique since the rotation around the vector $\boldsymbol{d}_{st}$ is not specified. Due to this, if all coefficients are included in the calculation of $\boldsymbol{F}_m$ the result will not be equivariant with respect to rotations, since their values will change based on the arbitrary rotation around $\boldsymbol{d}_{st}$. If rotation equivariance is desired for messages, the input and output coefficients to $\boldsymbol{F}_m$ can be restricted to only those for which $m = 0$. The $m = 0$ spherical harmonic functions are symmetric on the z-axis; see supplementary material and bases highlighted in green in Figure 1(a). Since the only remaining degree of freedom are rotations about the z-axis after rotating $\boldsymbol{x}_s^{(k)}$ and $\boldsymbol{x}_t^{(k)}$ by $\mathbf{A}_{st}$, the resulting $m = 0$ coefficients are invariant to rotations. Since the other inputs $|\boldsymbol{d}_{st}|$, a_s and a_t to $\boldsymbol{F}_m$ are also invariant to rotations, $\boldsymbol{F}_m$ is invariant to rotations if only $m = 0$ coefficients are used. The resulting messages m_{st} are equivariant to rotations once they are rotated back to the global coordinate frame using A_{st}^{-1}. While equivariance is a desirable property, as we demonstrate later, using $m \neq 0$ coefficients can help improve performance even though equivariance is violated.

The number of coefficients used during message passing when calculating $\boldsymbol{F}_m$ may be varied based on the edge properties to reduce the memory required by the network. For instance, atoms that are far away from each other may not need the same resolution of the spherical functions as those close together, so a lower maximum degree l may be used. By training several message passing networks with varying degrees for different ranges of edge distances, memory usage can be reduced while not impacting accuracy. Similarly, the range of m can typically be reduced to $[-1, 1]$ or $[-2, 2]$ since m for values greater than 2 (or less than -2) don't vary as much in the direction of the z-axis, reducing

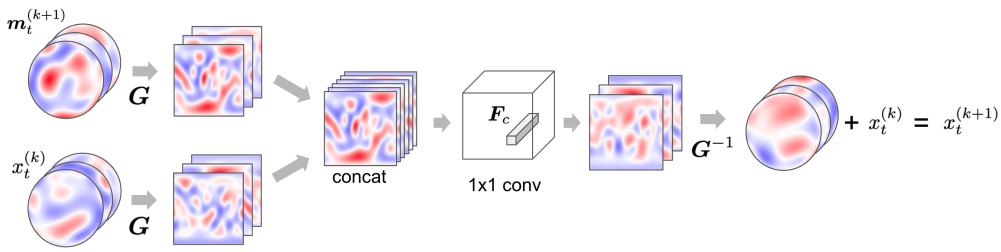

Figure 3: Illustration of message aggregation. The summed messages $\boldsymbol{m}_t^{(k+1)} = \sum_s \boldsymbol{m}_{st}^{(k+1)}$ and previous embedding $\boldsymbol{x}_t^{(k)}$ are converted from a spherical harmonic representation to a spherical grid representation using $\boldsymbol{G}$. The channels are concatenated and passed through a 3 layer 1×1 CNN. Each layer is followed by a SiLU activation. Finally, the channels are converted back to a spherical harmonic representation using $\boldsymbol{G}^{-1}$ and added to $\boldsymbol{x}_t^{(k)}$ to create $\boldsymbol{x}_t^{(k+1)}$.

their observed utility to the network. See Figure 1(b) for examples of the spherical channels when the range of m is reduced. Note that while the number of coefficients used by $\boldsymbol{F}_m$ may be reduced, the node embeddings maintain their original size based on the maximum L, so they may be able to aggregate information from all neighboring atoms at different distances and angles.

2.3 Message Aggregation

At this point in the network, the per edge messages are computed and rotated back to the global coordinate frame. To allow for complex interaction between the messages, we compute a non-linear function after summing all messages directed at a target node.

If we did this by naively passing the summed coefficients $\boldsymbol{m}_t^{(k+1)} = \sum_s \boldsymbol{m}_{st}^{(k+1)}$ through a fully connected network, the network may have difficulties learning representations that are approximately rotation equivariant (a fundamental property of atom forces). Another option is to place constraints on the non-linearities allowed to enforce equivariance using Clebsch–Gordan Nets [23] or steerable MLPs [4]. Instead, we apply an unconstrained pointwise function on the sphere, i.e., a function that is applied at every orientation without knowledge of the orientation from which the point was sampled [8]. For our function, we use a non-linear neural network $\boldsymbol{F}_c$ ($\mathbb{R}^C \rightarrow \mathbb{R}^C$) that combines information across channels. In practice, $\boldsymbol{F}_c$ is applied at a discrete number of orientations that results in a transformation that we numerically demonstrate (see supplementary material) is approximately equivariant to rotations:

$$\boldsymbol{x}_t^{(k+1)} = \boldsymbol{x}_t^{(k)} + \boldsymbol{G}^{-1}\left(\boldsymbol{F}_c\left(\boldsymbol{G}(\boldsymbol{m}_t^{(k+1)}), \boldsymbol{G}(\boldsymbol{x}_t^{(k)})\right)\right), \tag{3}$$

where $\boldsymbol{G}$ is a function that converts a spherical function represented by spherical harmonic coefficients, to one represented by point samples on a sphere. For $\boldsymbol{G}$ we use a spherical coordinate system (polar and azimuthal) to generate a 2D discrete representation of the spherical functions, see Figure 3. A 1×1 convolutional neural network $\boldsymbol{F}_c$ is applied to each discrete sample containing C channels, and the result is then converted back to the spherical harmonic representation using $\boldsymbol{G}^{-1}$. This is analogous to having a 2D image represented in the frequency domain, converting it to the spatial domain using an inverse discrete Fourier transform, applying a transformation and converting back to the frequency domain. Since the same operation is applied to all orientations on the sphere, a transformation that is approximately equivariant (up to discrete sampling limitations) to rotations may be learned, see the supplementary material for a more detailed discussion.

An illustration of the network is shown in Figure 3. To provide additional information to the network, $\boldsymbol{x}_t^{(k)}$ is also converted to a spherical grid representation, concatenated with $\boldsymbol{G}(\boldsymbol{m}_t^{(k+1)})$ and provided as input to the neural network to compute the final updates to the coefficients $\boldsymbol{x}_t^{(k+1)}$ at iteration $k+1$. Three layers of 1×1 convolutions on $2C$ channels are performed with each followed by a SiLU non-linear activation function. To avoid aliasing, the spherical grid is sampled at a resolution of $2 * (L+1)$.

Model	L	# layers	H	# batch	Samples / GPU sec.	OC20 S2EF Val Energy MAE [meV] ↓	Force MAE [meV/Å] ↓	Force Cos ↑	EFwT [%] ↑	OC20 IS2RE Val Energy MAE [meV] ↓	EwT [%] ↑
Median											
SchNet [34]						1400	78.3	0.109	0.00	-	-
DimeNet++ [20]						805	65.7	0.217	0.01	-	-
SpinConv [37]						406	36.2	0.479	0.13	-	-
GemNet-dT [22]					25.8	358	29.5	0.557	0.61	438	-
GemNet-OC [11]					18.3	**286**	25.7	0.598	1.06	407	-
SCN No rotation	4	12	512	64	8.5	410	67.7	0.192	0.01	-	-
SCN No 1x1 conv	6	12	1024	128	9.1	313	26.2	0.579	0.74	-	-
SCN $m = 0$	6	12	1024	128	10.4	307	26.5	0.588	0.83	-	-
SCN $m \in [-1, 1]$	6	12	1024	128	8.3	302	24.6	0.601	0.93	-	-
SCN $m \in [-2, 2]$	6	12	1024	96	6.9	301	23.4	0.612	1.01	-	-
SCN $m \in [-l, l]$	4	12	512	64	7.6	297	24.6	0.595	0.92	-	-
SCN	2	12	256	128	12.9	312	27.6	0.568	0.70	-	-
SCN	4	12	512	128	12.1	303	25.3	0.592	0.87	-	-
SCN	6	12	1024	128	8.3	302	24.6	0.601	0.93	-	-
SCN	8	12	1024	64	5.9	300	**23.2**	0.620	1.15	-	-
SCN $m = 0$	6	16	1024	96	7.4	300	25.7	0.600	0.96	394	9.63
SCN $m = 0$	8	16	1024	64	4.8	296	25.3	0.608	1.01	389	9.85
SCN	6	16	1024	96	5.9	**287**	**22.8**	**0.623**	**1.22**	371	10.47
SCN	8	16	1024	96	3.5	**283**	**22.7**	**0.627**	**1.22**	**364**	**11.26**

Table 1: Results on the OC20 2M training dataset and ablation studies for SCN model variations. The validation results are averaged across the four OC20 Validation set splits. All SCN models are trained on 16 GPUs for 12 epochs with the learning rate reduced by 0.3 at 5, 7, 9, and 11 epochs, except SCN with $L = 8$ and 16 layers that used 32 GPUs to obtain a larger batch size. Batch sizes vary based on the number of instances that can be fit in 32GB RAM.

2.4 Energy and Forces Estimation

We compute the energy E by first estimating a per-atom energy using a pointwise function on the sphere and taking its integral over all possible orientations. The per-atom energy estimates are summed to obtain the overall system's energy:

$$E = \sum_i \int_{\hat{\boldsymbol{r}}} \boldsymbol{F}_{energy} \left(s_i^{(K)}(\hat{\boldsymbol{r}}) \right), \tag{4}$$

where $\boldsymbol{F}_{energy}$ is a three layer fully connected neural network ($\mathbb{R}^C \to \mathbb{R}^1$) with SiLU activation functions. We calculate forces similarly by estimating a force magnitude $|\boldsymbol{f}| = \boldsymbol{F}_{force}(s_i^{(K)}(\hat{\boldsymbol{r}}))$ in every direction $\hat{\boldsymbol{r}}$ over the sphere. Integration is performed after multiplying the magnitude by the orientation $\hat{\boldsymbol{r}}$ to obtain directional vectors:

$$\boldsymbol{f}_i = \int_{\hat{\boldsymbol{r}}} \hat{\boldsymbol{r}} \boldsymbol{F}_{force} \left(s_i^{(K)}(\hat{\boldsymbol{r}}) \right), \tag{5}$$

where $\boldsymbol{F}_{force}$ is a three layer fully connected neural network ($\mathbb{R}^C \to \mathbb{R}^1$). In practice, a discrete approximation of the integral is performed for equations (4) and (5) using a set of 128 evenly distributed points on the sphere. Note that the application of a pointwise function of an arbitrary neural network at a finite number of discrete orientations is not invariant (energies) or equivariant (forces) to rotations. However in practice, numerically this operation does closely approximate these properties (see Supplementary B.2).

3 Experiments

We present results on the Open Catalyst 2020 (OC20) dataset [5] that is released under a Creative Commons Attribution 4.0 License. OC20 contains over 130M training examples for the task of predicting atomic energies and forces for catalysts used in renewable energy storage and other important applications [44]. This dataset is a popular benchmark for the ML community. We begin by comparing results across all tasks on the test set. Since several papers only report results on the IS2RE task using approaches that directly predict relaxed energies, we also train a model specifically for this task and approach, and compare. Finally, we show numerous ablation studies comparing model variations on the smaller OC20 2M dataset.

Model	#Params	**Train time**	S2EF Test Energy MAE meV ↓	Force MAE [meV/Å] ↓	Force Cos ↑	EFwT [%] ↑	IS2RE Test Energy MAE meV ↓
Median	–		2258	84.4	0.016	0.01	
SchNet [34, 5]	9.1M	194d	540	54.7	0.302	764	
PaiNN [32]	20.1M	67d	341	33.1	0.491	0.46	471
DimeNet++-L-F+E [20, 5]	10.7M	1600d	480	31.3	0.544	0.00	559
SpinConv (force-centric) [37]	8.5M	275d	336	29.7	0.539	0.45	437
GemNet-dT [22]	32M	492d	292	24.2	0.616	1.20	400
GemNet-OC [11]	39M	336d	**233**	20.7	0.666	2.50	355
GemNet-OC-L-E + MD [11]	56M	640d	**230**	21.0	0.665	2.80	-
GemNet-OC-L-F + MD [11]	216M	765d	241	19.0	**0.691**	**2.97**	-
GemNet-OC-L-F+E + MD [11]	-	-	-	-	-	-	348
SCN + MD $m = 0$	122M	784d	248	19.7	0.685	2.33	350
SCN + MD	271M	832d	238	**17.5**	**0.691**	2.65	**322**

Table 2: Comparison of SCN to existing GNN models on the S2EF and IS2RE tasks trained on the All and MD datasets. Average results across all four test splits are reported. We mark as bold the best performance and close ones, *i.e.*, within 0.5 meV/Å MAE, which we found to empirically provide a meaningful performance difference. Training time is in GPU days. Median represents the trivial baseline of always predicting the median training energy or force across all the validation atoms. The SCN model has $L = 8$, 20 layers, $C = 128$ and $E = 256$. See supplementary material for SCN results trained on the OC20 All dataset and IS2RS results.

3.1 Implementation Details

The implementation of spherical harmonics and their transformations uses the code provided by Euclidean neural networks (e3nn) [12]. For message passing, two resolutions of spherical harmonics are used based on a atom's nearest neighbors, see the supplementary material (Section A) for exact parameters. Unless otherwise stated $C = 128$, $K = 16$, $H = 1024$, $E = 128$, and for $\boldsymbol{F}_m$ in message passing only $m \in [-1, 1]$. Training is performed using the AdamW optimizer [26] with a learning rate of 0.0004. The effective batch size is increased using data parallelism and NVIDIA's Automatic Mixed Precision (AMP). All model code will be open sourced with an MIT license upon acceptance. Please see the supplementary for details on the PaiNN[32] and NequIP[2] baselines.

3.2 OC20 Tasks

We compare against numerous models trained on the OC20 2M, All and MD datasets across various tasks [5]. OC20 2M contains 2M, OC20 All has 133M and OC20 MD has 38M training examples. The 16 layer SCN models significantly outperform state-of-the-art GemNet-OC [11] on S2EF force ($\approx 11\%$ improvement) and IS2RE energy predictions ($\approx 10\%$ improvement) when trained on OC20 2M. Comparable results are found for energy MAE on S2EF with GemNet-OC. We also report throughout efficiency. Note that while the SCN models process fewer samples per second, they are much more sample efficient as shown in Figure 2(b).

We compare our largest SCN models trained on the All+MD datasets (133M + 38M examples) across all OC20 tasks in Table 2. Our model achieves state-of-the-art results for force MAE ($\approx 8\%$ improvement) and force cosine for the S2EF tasks. Only the recently released GemNet-OC model [11] performs better on energy MAE for S2EF. On the challenging IS2RE task that requires models to predict both accurate forces and energies, SCN outperforms GemNet-OC by over 7%. Note that we only train a single model to predict both energy and forces, where GemNet-OC-L-F+E + MD is the combination of two models, one used for forces and the other for energy. Please see the supplementary material for results trained on the All dataset and IS2RS results.

3.3 OC20 IS2RE Direct

The task of Initial Structure to Relaxed Energy (IS2RE) may be accomplished using two approaches: 1) Use an S2EF model to relax the positions of the atoms (find local energy minima) and output the energy at the relaxed structure, or 2) directly predict the relaxed energy from the initial structure

OC20 IS2RE Direct Test

Model	Energy MAE [meV] ↓					EwT [%] ↑			
	ID	OOD Ads	OOD Cat	OOD Both	Average	ID	OOD Ads	OOD Cat	OOD Both
Median baseline	1750	1879	1709	1664	0.71	0.72	0.89	0.74	
CGCNN[41]	615	916	622	851	751	3.40	1.93	3.10	2.00
SchNet[34]	639	734	662	704	685	2.96	2.33	2.94	2.21
NequIP ($l = 2$)[2]	602	784	619	736	685	3.18	2.16	3.06	2.12
PaiNN [32]	575	783	604	743	676	3.46	1.97	3.46	2.28
GemNet-dT[22]	527	758	549	702	634	4.59	2.09	4.47	2.28
DimeNet++[20]	562	725	576	661	631	4.25	2.07	4.10	2.41
GemNet-OC[11]	560	711	576	671	630	4.15	2.29	3.85	2.28
SphereNet[25]	563	703	571	638	619	4.47	2.29	4.09	2.41
SEGNN[4]	533	692	537	679	610	**5.37**	2.46	**4.91**	2.63
SCN	**516**	**643**	**530**	**604**	**573**	4.92	**2.71**	4.42	**2.76**

Table 3: Results on IS2RE OC20 Test for approaches that directly predict the relaxed energies without performing relaxations and do not use auxiliary losses during training. See Tables 1 and 2 for results using relaxation based approaches to IS2RE. Results are shown for two metrics; energy Mean Absolute Error (MAE) and Energy within Threshold (EwT). The SCN model uses $L = 6$ with 16 layers and is trained for 21 epochs.

without performing the relaxation. Empirically, the first approach achieves higher accuracy (see Tables 1 and 2), while the second approach is more efficient both during training and inference. While solving the same problem, the second approach of direct prediction is a fundamentally different problem than the first. The direct approach cannot take advantage of detailed position and angular information, since the initial structure only contains a rough placement of the atoms. Instead it must rely more on global knowledge of how the atoms are arranged. Thus, an approach that works well for relaxation based approaches to IS2RE might not work well for direct approaches.

As shown in Table 3, the SCN model achieves state-of-the-art results for energy MAE across all test splits for models that do not use ancillary loses [42, 14]. Energy within Threshold (EwT) is comparable to SEGNN [4]. Improvements for SCN are most pronounced for the out-of-domain splits, demonstrating its increased ability to generalize to unseen data.

3.4 Ablation studies

We explore numerous model variations in Table 1 trained on the OC20 2M dataset and evaluated on the OC20 validation set. The first set of ablation experiments explores reducing the complexity of our model. If messages are not rotated before and after applying $\boldsymbol{F}_m$ in Equation 2, *i.e.*, $\mathbf{A}_{st}$ is an identity matrix, the results are significantly worse for energies and even more so for forces since forces are highly dependent on angular information. If we replace the non-linear aggregation (Equation 3) with a simple summation, the results also degrade across all metrics. Results improve for $m \in [-1, 1]$ as compared to only using $m = 0$ coefficients, which demonstrates that if the equivariant to rotation constraint is relaxed, improved results may be achieved. For higher m, $m \in [-2, 2]$ or $m \in [-4, 4]$ diminishing returns are noticed with increased computational cost.

When scaling the network, increasing L has a more significant impact on force estimation, while depth improves energy estimation. This may be due to higher L resulting in greater angular resolution, which is important to force estimation. Greater depth, which allows information to be shared more often between atoms, may lead to better energy predictions since energy is function of the entire atomic structure.

4 Related work

Molecular and atomic property prediction has made significant recent progress. Models relying on hand-crafted representations like Behler-Parrinello[3] and sGDML[7] have recently been surpassed by learned feature representations using GNNs [2, 20, 22, 32]. Early GNN developments focused only on invariant representations. CGCNN[41] and SchNet[31] make energy and force estimations using only atomic species and pair-wise distance information. DimeNet[20, 21], SphereNet[25], and

GemNet[22, 11] extend this to explicitly capture triplet and quadruplet angles, which are scalars. Utilizing invariant representations offers flexibility in model architecture design since the node or edge features are inherently invariant and model constraints are not needed to maintain equivariance. The enumeration of triplets and quadruplets of atoms to calculate functions of relative or dihedral angles requires careful model design to maintain efficiency. More recently, equivariant models, capturing both scalar and equivariant features, have outperformed traditional GNNs on small molecular datasets including MD17[6] and QM9[28]. While not strictly equivariant, our work falls into this later class of GNNs for atomic property predictions.

Models that strictly enforce equivariance to the SO(3) (3D rotations) or E(3) group (3D rigid transformations) share similarities to our use of spherical channels [32, 1, 2, 8, 4]. Both represent atoms by higher-order tensors and not just scalar information. For instance, Tensor Field Networks [38], Cormorant [1], NequIP [2] and SEGNN [4] use the same spherical harmonic representations as SCNs, while PaiNN [32] is an $l = 1$ variation. All models place constraints on the operations that can be performed to ensure equivarance, such as restricting non-linear operators to invariant scalar inputs, e.g., pair-wise atom distances. In networks like TFN, geometric information of higher degree l are mixed into the $l = 0$ features via tensor products on which a gated non-linear function is applied. Similar to our approach, the recently proposed SEGNN uses the spherical harmonic representation to enable steerable features based on edge orientations. SEGNN introduces steerable MLPs that enforce equivariance by restricting the learnable parameters to those that scale the Clebsch-Gordan tensor products followed by gated non-linearities [39]. In our approach, we increase the expressivity of the model by identifying a set of $m = 0$ coefficients across degrees l that are invariant to rotations given an edge's orientation to which unconstrained functions may be applied. We also demonstrate that improved performance can be achieved by relaxing the requirement for strict message equivariance by including a broader set of coefficients beyond $m = 0$. Furthermore, we project the spherical channels onto a grid and perform a pointwise 1×1 convolution followed by a non-linearity. Although not strictly equivariant, the pointwise operation allows for complex mixing of different degrees of spherical harmonics resulting in rich geometric descriptions.

5 Discussion

A limitation of SCN is that it directly predicts forces, so energy conservation is not enforced. This may limit its use in chemistry applications such as molecular dynamics. Non-energy conserving models are still practically useful for applications such as structure relaxations and transition state searches. The SCN model scales $O(n^2)$ with respect to L, which may limit using larger L. Using a larger set of m coefficients during message passing also faces challenges, since the network may find it more difficult to learn approximately equivariant mappings. In Table 1, we saw that increased depth improves results, especially to energy prediction. It remains an open question whether deeper networks will see further improvement.

While this work is motivated by problems we face in addressing climate change, advances in chemistry may have numerous use cases. Some of these use cases may have unintended consequences, such as the development of the Haber-Bosch process for ammonia fertilizer production that enabled the world to feed a growing population, but over fertilization has recently led to ocean "dead zones" and its production is very carbon intensive. More alarmingly, the knowledge gained from fertilizer production was used to create explosives during wartime [16]. We hope to steer research in this area towards beneficial applications by utilizing datasets such as OC20.

In conclusion, we propose a GNN that uses spherical harmonics to represent channel embeddings that contain explicit orientation information. While the network is encouraged to learn output mappings that are equivariant to rotations, best performance is achieved by relaxing the equivariant constraints and allowing for more expressive non-linear transformations. We demonstrate state-of-the-art results across numerous tasks on the large-scale OC20 dataset for modeling catalyst for applications addressing climate change.

References

[1] Brandon Anderson, Truong-Son Hy, and Risi Kondor. Cormorant: Covariant molecular neural networks. *arXiv preprint arXiv:1906.04015*, 2019.

[2] Simon Batzner, Tess E Smidt, Lixin Sun, Jonathan P Mailoa, Mordechai Kornbluth, Nicola Molinari, and Boris Kozinsky. Se (3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *arXiv preprint arXiv:2101.03164*, 2021.

[3] Jörg Behler. Perspective: Machine learning potentials for atomistic simulations. *The Journal of chemical physics*, 145(17):170901, 2016.

[4] Johannes Brandstetter, Rob Hesselink, Elise van der Pol, Erik Bekkers, and Max Welling. Geometric and physical quantities improve e (3) equivariant message passing. *arXiv preprint arXiv:2110.02905*, 2021.

[5] Lowik Chanussot, Abhishek Das, Siddharth Goyal, Thibaut Lavril, Muhammed Shuaibi, Morgane Riviere, Kevin Tran, Javier Heras-Domingo, Caleb Ho, Weihua Hu, Aini Palizhati, Anuroop Sriram, Brandon Wood, Junwoong Yoon, Devi Parikh, C. Lawrence Zitnick, and Zachary Ulissi. Open catalyst 2020 (oc20) dataset and community challenges. *ACS Catalysis*, 11(10):6059–6072, 2021.

[6] 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.

[7] Stefan Chmiela, Huziel E Sauceda, Klaus-Robert Müller, and Alexandre Tkatchenko. Towards exact molecular dynamics simulations with machine-learned force fields. *Nature communications*, 9(1):1–10, 2018.

[8] Taco S Cohen, Mario Geiger, Jonas Köhler, and Max Welling. Spherical cnns. In *ICLR*, 2018.

[9] Carlos Esteves, Christine Allen-Blanchette, Ameesh Makadia, and Kostas Daniilidis. Learning so (3) equivariant representations with spherical cnns. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 52–68, 2018.

[10] William T Freeman, Edward H Adelson, et al. The design and use of steerable filters. *IEEE Transactions on Pattern analysis and machine intelligence*, 13(9):891–906, 1991.

[11] Johannes Gasteiger, Muhammed Shuaibi, Anuroop Sriram, Stephan Günnemann, Zachary Ulissi, C Lawrence Zitnick, and Abhishek Das. How do graph networks generalize to large and diverse molecular systems? *arXiv preprint arXiv:2204.02782*, 2022.

[12] Mario Geiger, Tess Smidt, Alby M., Benjamin Kurt Miller, Wouter Boomsma, Bradley Dice, Kostiantyn Lapchevskyi, Maurice Weiler, Michał Tyszkiewicz, Simon Batzner, Martin Uhrin, Jes Frellsen, Nuri Jung, Sophia Sanborn, Josh Rackers, and Michael Bailey. Euclidean neural networks: e3nn, 2020. URL `https://doi.org/10.5281/zenodo.5292912`.

[13] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International Conference on Machine Learning (ICML)*, pages 1273–1272, 2017.

[14] Jonathan Godwin, Michael Schaarschmidt, Alexander L Gaunt, Alvaro Sanchez-Gonzalez, Yulia Rubanova, Petar Veličković, James Kirkpatrick, and Peter Battaglia. Simple gnn regularisation for 3d molecular property prediction and beyond. In *International Conference on Learning Representations*, 2021.

[15] Marco Gori, Gabriele Monfardini, and Franco Scarselli. A new model for learning in graph domains. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, volume 2, pages 729–734. IEEE, 2005.

[16] Thomas Hager. *The alchemy of air: a Jewish genius, a doomed tycoon, and the scientific discovery that fed the world but fueled the rise of Hitler*. Broadway Books, 2009.

[17] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016.

[18] Weihua Hu, Muhammed Shuaibi, Abhishek Das, Siddharth Goyal, Anuroop Sriram, Jure Leskovec, Devi Parikh, and C Lawrence Zitnick. Forcenet: A graph neural network for large-scale quantum calculations. *arXiv preprint arXiv:2103.01436*, 2021.

[19] Peter Bjørn Jørgensen, Karsten Wedel Jacobsen, and Mikkel N Schmidt. Neural message passing with edge updates for predicting properties of molecules and materials. *arXiv preprint arXiv:1806.03146*, 2018.

[20] Johannes Klicpera, Shankari Giri, Johannes T. Margraf, and Stephan Günnemann. Fast and uncertainty-aware directional message passing for non-equilibrium molecules. In *NeurIPS-W*, 2020.

[21] Johannes Klicpera, Janek Groß, and Stephan Günnemann. Directional message passing for molecular graphs. In *International Conference on Learning Representations (ICLR)*, 2020.

[22] Johannes Klicpera, Florian Becker, and Stephan Günnemann. Gemnet: Universal directional graph neural networks for molecules. *arXiv preprint arXiv:2106.08903*, 2021.

[23] Risi Kondor, Zhen Lin, and Shubhendu Trivedi. Clebsch–gordan nets: a fully fourier space spherical convolutional neural network. *Advances in Neural Information Processing Systems*, 31, 2018.

[24] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

[25] Yi Liu, Limei Wang, Meng Liu, Xuan Zhang, Bora Oztekin, and Shuiwang Ji. Spherical message passing for 3d graph networks. *arXiv preprint arXiv:2102.05013*, 2021.

[26] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019.

[27] Zhuoran Qiao, Matthew Welborn, Animashree Anandkumar, Frederick R Manby, and Thomas F Miller III. Orbnet: Deep learning for quantum chemistry using symmetry-adapted atomic-orbital features. *arXiv preprint arXiv:2007.08026*, 2020.

[28] 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.

[29] David Rolnick, Priya L Donti, Lynn H Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran, Andrew Slavin Ross, Nikola Milojevic-Dupont, Natasha Jaques, Anna Waldman-Brown, et al. Tackling climate change with machine learning. *arXiv preprint arXiv:1906.05433*, 2019.

[30] Victor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E (n) equivariant graph neural networks. *arXiv preprint arXiv:2102.09844*, 2021.

[31] Kristof Schütt, Pieter-Jan Kindermans, Huziel Enoc Sauceda Felix, Stefan Chmiela, Alexandre Tkatchenko, and Klaus-Robert Müller. Schnet: A continuous-filter convolutional neural network for modeling quantum interactions. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 991–1001, 2017.

[32] Kristof Schütt, Oliver Unke, and Michael Gastegger. Equivariant message passing for the prediction of tensorial properties and molecular spectra. In *International Conference on Machine Learning*, pages 9377–9388. PMLR, 2021.

[33] Kristof T Schütt, Farhad Arbabzadah, Stefan Chmiela, Klaus R Müller, and Alexandre Tkatchenko. Quantum-chemical insights from deep tensor neural networks. *Nature communications*, 8:13890, 2017.

[34] Kristof T Schütt, Huziel E Sauceda, P-J Kindermans, Alexandre Tkatchenko, and K-R Müller. Schnet–a deep learning architecture for molecules and materials. *The Journal of Chemical Physics*, 148(24):241722, 2018.

[35] Andrew W Senior, Richard Evans, John Jumper, James Kirkpatrick, Laurent Sifre, Tim Green, Chongli Qin, Augustin Žídek, Alexander WR Nelson, Alex Bridgland, et al. Improved protein structure prediction using potentials from deep learning. *Nature*, 577(7792):706–710, 2020.

[36] Connor Shorten and Taghi M Khoshgoftaar. A survey on image data augmentation for deep learning. *Journal of big data*, 6(1):1–48, 2019.

[37] Muhammed Shuaibi, Adeesh Kolluru, Abhishek Das, Aditya Grover, Anuroop Sriram, Zachary Ulissi, and C. Lawrence Zitnick. Rotation invariant graph neural networks using spin convolutions. *arXiv preprint arXiv:2106.09575*, 2021.

[38] Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds. *arXiv preprint arXiv:1802.08219*, 2018.

[39] Maurice Weiler, Mario Geiger, Max Welling, Wouter Boomsma, and Taco Cohen. 3d steerable cnns: Learning rotationally equivariant features in volumetric data. *arXiv preprint arXiv:1807.02547*, 2018.

[40] Maurice Weiler, Fred A Hamprecht, and Martin Storath. Learning steerable filters for rotation equivariant cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 849–858, 2018.

[41] Tian Xie and Jeffrey C Grossman. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Physical review letters*, 120(14):145301, 2018.

[42] Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and Tie-Yan Liu. Do transformers really perform badly for graph representation? *Advances in Neural Information Processing Systems*, 34, 2021.

[43] Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applications. *AI Open*, 1:57–81, 2020.

[44] C. L. Zitnick, L. Chanussot, A. Das, S. Goyal, J. Heras-Domingo, C. Ho, W. Hu, T. Lavril, A. Palizhati, M. Riviere, M. Shuaibi, A. Sriram, K. Tran, B. Wood, J. Yoon, D. Parikh, and Z. Ulissi. An introduction to electrocatalyst design using machine learning for renewable energy storage. *arXiv preprint arXiv:2010.09435*, 2020.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] , see section 5.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] , see section 5.

(d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]

2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No] , code will be released with MIT license upon acceptance or sooner[N/A] .
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] , experiments are too expensive to run multiple times. However, we provide rough estimates for run variations in Table 1.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] , training times and throughput are proved.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]