A universal framework for featurization of atomistic systems

Xiangyun Lei

Georgia Institute of Technology

Andrew Medford (✉ andrew.medford@chbe.gatech.edu)

Georgia Institute of Technology <https://orcid.org/0000-0001-8311-9581>

Article

Keywords: Gaussian multipole, neural networks, molecular dynamics

Posted Date: October 5th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-952157/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

A UNIVERSAL FRAMEWORK FOR FEATURIZATION OF ATOMISTIC SYSTEMS

A PREPRINT

Xiangyun Lei

School of Chemical and Biomolecular Engineering
Georgia Institute of Technology
Atlanta, GA, 30318 USA
xlei38@gatech.edu

Andrew J. Medford

School of Chemical and Biomolecular Engineering
Georgia Institute of Technology
Atlanta, GA, 30318 USA
ajm@gatech.edu

October 1, 2021

ABSTRACT

Molecular dynamics simulations are an invaluable tool in numerous scientific fields. However, the ubiquitous classical force fields cannot describe reactive systems, and quantum molecular dynamics are too computationally demanding to treat large systems or long timescales. Reactive force fields based on physics or machine learning can be used to bridge the gap in time and length scales, but these force fields require substantial effort to construct and are highly specific to a given chemical composition and application. A significant limitation of machine learning models is the use of element-specific features, leading to models that scale poorly with the number of elements. This work introduces the Gaussian multipole (GMP) featurization scheme that utilizes physically-relevant multipole expansions of the electron density around atoms to yield feature vectors that interpolate between element types and have a fixed dimension regardless of the number of elements present. We combine GMP with neural networks to directly compare it to the widely used Behler-Parinello symmetry functions for the MD17 dataset, revealing that it exhibits improved accuracy and computational efficiency. Further, we demonstrate that GMP-based models can achieve chemical accuracy for the QM9 dataset, and their accuracy remains reasonable even when extrapolating to new elements. Finally, we test GMP-based models for the Open Catalysis Project (OCP) dataset, revealing comparable performance to graph convolutional deep learning models. The results indicate that this featurization scheme fills a critical gap in the construction of efficient and transferable machine-learned force fields.

Atomistic simulations are a crucial tool in many scientific fields, ranging from protein engineering to materials design [1, 2, 3, 4, 5, 6, 7, 8, 9]. Full quantum mechanical treatment of atoms provides highly accurate energies and forces, but the computational cost is prohibitive for the length and time scales relevant to most applications [10, 11, 12]. Classical and reactive force fields can act as surrogates for the quantum mechanical simulations, enabling simulations at longer length and time scales [13, 14, 15, 16, 17, 18, 7, 19]. However, these models are specialized to specific systems and have a limited ability to simulate inherently quantum-mechanical phenomena such as covalent bond formation [20]. Machine-learning models have recently emerged as a promising strategy to fill the gap between quantum mechanical simulations and classical force field models [21, 22, 23]. The field of machine-learned force fields has exploded in the last decade, leading to a plethora of different machine-learning force field models capable of predicting energies and forces with accuracy comparable to the underlying method [24, 25, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35, 36, 37, 38]. However, most models are customized for specific application domains, and a framework for a general-purpose model has not been established.

One fundamental problem in machine-learning models for atomistic systems is the issue of feature generation. Cartesian coordinates and elemental identities are the most common way to define atomistic systems. However, this description does not capture the system’s rotational, translational, or permutation invariances. Converting the Cartesian coordinates to “feature vectors” that describe each atom is commonly used to solve this problem. Researchers have devised a wide range of featurization strategies that encode these fundamental physical symmetries, recently classified and thoroughly covered in several thorough reviews [39, 40]. Some examples include the Coulomb matrix[41], radial

distribution function-based fingerprint [42], FCHL[35], SOAP [36, 37], COMB [17], Chebyshev polynomials [43] and the atom-centered symmetry functions[24, 17]. These strategies have yielded remarkably accurate models within various sub-fields but are not sufficiently general to treat all atomistic systems.

The most common limitation is related to the scaling of the feature vector size as the number of elements in the system increases. Most existing descriptors scale polynomially or combinatorially with the number of elements in the system [24, 36]. This poor scaling means that these featurization strategies can only be applied to training sets with a limited number of elements, making the prospect of a universal machine-learning model that works for all elements infeasible. While some approaches overcome this limitation [38, 43], they have not been tested in on large many-element benchmark systems. In addition to fingerprint scalability, there is also the issue of scalability of the regression model. Many high-accuracy models rely on kernel ridge regression (KRR), a technique that suffers from poor scalability at with extremely large training set sizes. Neural networks are far more scalable, but if typical element-specific featurization schemes are used then a high-dimensional neural network (HDNN) [24, 44] is required to connect the features to atomic properties. The HDNN scheme uses element-specific neural networks, so that the number of model parameters also increases with the number of elements. An alternative to featurization is deep learning [45], typically with graph-convolutional neural network (GCN) models [25, 26, 27, 28, 32, 29, 30]. These deep learning models show an excellent ability to learn appropriate representations for molecular, solid-state, and surface systems with many elements [46]. However, this comes at the cost of less transparent models that typically require more time for training and prediction than their feature-based counterparts. Moreover, many of the deep learning approaches utilize elemental features [32, 25, 27, 28], suggesting that improved featurization schemes will translate to improved deep learning models.

Here we introduce a new featurization approach called Gaussian multipole (GMP) features. The GMP features utilize an implicit description of the electron density so that the feature vector’s dimension is independent of the number of elements. Using the electron density as the fundamental input makes them suitable for universal machine-learning models that work for all elements and provides a straightforward route for extending them to systems that involve charged atoms or magnetic moments. It also naturally allows the use of the more efficient single neural network (SNN) [33] structure, where all elements share the same neural network. Moreover, the GMP features are related to a multipole expansion of the electron density [47, 48], making them physically relevant and systematically improvable. Remarkably, we show that the GMP features are capable of interpolating between elemental species, a capability that to our knowledge has not been demonstrated with any other featurization schemes. These properties of the GMP features will facilitate a new family of fast and interpretable feature-based machine learning models that have the general applicability of more complex and opaque GCN models.

The GMP approach encodes the elemental identity through an approximate description of the electron density based on Gaussian basis functions. In this work, we use the valence density extracted from the SG15 pseudopotentials [49] approximated by 2-7 atom-centered Gaussians per element, normalized by the total number of valence electrons. The number and widths of the Gaussians for each element are determined using nonlinear regression, and these “static valence densities” are fixed for all models presented in this work. Details are provided in the supporting information. Conceptually, this is equivalent to using the valence density of non-interacting atoms as the fundamental description of the system (Fig. 1a). This approach also allows for interpolation and extrapolation between different elemental species, and concepts like charged atoms can be naturally accommodated using the isolated charged atom to construct a suitable valence potential. In principle, it is also possible to use other quantities like the all-electron density, spin density, or self-consistent electron density, as long as the quantity can be represented by a linear combination of atom-centered Gaussians.

The approximated electron density is then vectorized via the inner product between the electron density and atom-centered “probe” functions to generate the features for the atoms. The probe function consists of a Gaussian function and a Maxwell-Cartesian spherical harmonic (MCSH) function. Gaussian functions of varying width account for the radial variation (Fig. 1b) of surrounding electron density of an atom, and the MCSH functions capture angular variation (Fig. 1c). The Gaussian product rule provides analytical solutions to the integral in Eqn. 1, enabling highly efficient computation of the features. Rotational invariance is enforced by taking norms of each MCSH group [50, 51].

Figure 1: Illustration for the GMP featurization scheme. a) Construct the electron density distribution using linear combination of Gaussians. b) Apply the Gaussian radial probe to focus on the electronic environments at different radial length scales around each atom core by masking the electron density distribution with Gaussian functions of different widths. c) Apply the angular (MCSH) probe that acts as a multi-pole expansion of the radially-masked electron density via inner products with different groups of MCSH functions. d) Take the norm of the results for each group to ensure rotational invariance, yielding an entry to the feature vector for each individual atom.

Mathematically this is described as

$$\begin{aligned}
 \mu_{i,abc} &= \langle probe, \hat{\rho} \rangle = \langle angular\ probe \times radial\ probe, \hat{\rho} \rangle \\
 &= \langle S_{abc} \times G_i, \sum_j \sum_k G_{dens,j,k} \rangle \\
 &= \iiint_V S_{abc} G_{probe,i} \sum_j \sum_k G_{dens,j,k} dV \\
 &= \sum_j \sum_k \iiint_V S_{abc} G_{probe,i} G_{dens,j,k} dV
 \end{aligned} \tag{1}$$

where $\langle a, b \rangle$ denotes inner product of two functions, V is the volume, $\mu_{i,abc}$ is the feature resulting from the radial probe G_i and angular probe S_{abc} . $\hat{\rho}$ is the distribution of electron density of a molecule, approximated by linear combinations of primitive Gaussians $G_{dens,j,k}$ centered at each atom. The formalism does not include a strict cutoff radius, though for computational efficiency the sum can be restricted to surrounding atoms with a non-negligible overlap. The MCSH functions S_{abc} are linear combinations of polynomials:

$$S_{abc}^n = \sum_{\substack{t \\ \text{terms}}} C_t x^{m_{x,t}} y^{m_{y,t}} z^{m_{z,t}} \quad (2)$$

where abc is the specific index of the spherical harmonic function, $n = a + b + c$ is the order of it, and m_x, m_y, m_z are the exponents of the specific polynomial. The first 4 orders of MCSH functions are listed in Table 1.

n	group	{abc}	$S_{abc}^{(n)}$	n	group	{abc}	$S_{abc}^{(n)}$
0	1	000	1	3	1	300	$15x^3 - 9x$
1	1	100	x			030	$15y^3 - 9y$
		010	y			003	$15z^3 - 9z$
		001	z	2	210	$15x^2y - 3y$	
2	1	200	$3x^2 - 1$		201	$15x^2z - 3z$	
		020	$3y^2 - 1$		021	$15y^2z - 3z$	
		002	$3z^2 - 1$		120	$15xy^2 - 3x$	
	2	110	$3xy$		102	$15xz^2 - 3x$	
		101	$3xz$		012	$15yz^2 - 3y$	
		011	$3yz$	3	111	$15xyz$	

Table 1: The analytical expressions of the first four orders of MCSH denoted by $S_{abc}^{(n)}$

To ensure that the resulting features are rotationally invariant, we use the norm of each group, $\mu_{i,abc}$ [50]. Therefore, the GMP feature vector is defined as

$$\vec{\Phi} = \Phi_{i,abc} = \sqrt{\sum_{P(a,b,c)} \mu_{i,abc}^2} \mid a, b, c \in \mathbb{N}, \quad (3)$$

where Φ denotes the GMP features, i is an index over the radial probes, abc is an index combination corresponding to a rotational group, and $P(a, b, c)$ denotes the permutation group of a, b, c (e.g. $P(1, 0, 0) = \{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$). The set of features can thus be written as

$$\vec{\Phi} = \left\{ \begin{array}{l} \sqrt{\mu_{1,000}^2}, \sqrt{\mu_{1,100}^2 + \mu_{1,010}^2 + \mu_{1,001}^2}, \sqrt{\mu_{1,200}^2 + \mu_{1,020}^2 + \mu_{1,002}^2}, \sqrt{\mu_{1,110}^2 + \mu_{1,101}^2 + \mu_{1,011}^2}, \dots \\ \sqrt{\mu_{2,000}^2}, \sqrt{\mu_{2,100}^2 + \mu_{2,010}^2 + \mu_{2,001}^2}, \sqrt{\mu_{2,200}^2 + \mu_{2,020}^2 + \mu_{2,002}^2}, \sqrt{\mu_{2,110}^2 + \mu_{2,101}^2 + \mu_{2,011}^2}, \dots \\ \dots \end{array} \right\}. \quad (4)$$

Conceptually, the resulting features are similar to a multipole expansion of the electron density around each atom. The multipole expansion provides rotational features that are complete and orthogonal. The Gaussian probe’s width controls the radial length scale of the multi-pole expansions, and the radial features are over-complete. These properties reduce linear dependencies within the features and lead to a systematic improvement in the system’s description as the number of features increases. The GMP scheme is closely related to the “wavelet” features [38], but it uses a different representation of the electron density, and differs in the way the spherical harmonics are normalized and computed, so that it provides atom-centered features. It is also similar to the SOAP scheme, which uses element-specific densities and spherical harmonics [24, 36, 37]. However, the key difference is that the length of the feature vector is fixed regardless of the number of elements, and the normalization across rotational groups leads to fewer features. These traits that enable GMP to efficiently featurize complex systems with an arbitrary number of elements and arbitrary boundary conditions.

In this work, we combine the GMP features with a neural network regression model to demonstrate the efficiency, accuracy, and transferability of models based on these features. As mentioned, the single neural network (SNN) architecture [33] is well-suited to the GMP features since all elements utilize the same features (Fig. 1). We also utilize a per-element bias term, equivalent to fitting to formation energies instead of total energies, to reduce the magnitude

of the energies. We have implemented the GMP+SNN framework in the AMPTorch code [52, 53] and used this implementation for all models in this work. The number of parameters used by the GMP+SNN model varies depending on the number of features, hidden layers, and nodes per layer but is generally lower than the number of parameters in a comparable GCN by an order of magnitude. Details of model fitting, implementation, and parameters are provided in the supplementary information. The notation $\text{GMP}(N_{\text{MCSH}}, N_{\text{Gaussian}})$ is used to denote the feature sets in the examples below, where N_{Gaussian} is the number of radial Gaussians, and N_{MCSH} is the maximum order MCSH used to construct the feature set. Note that although there are multiple groups in each order of MCSH, with different importance in describing the environment, we always include all the groups in each order in this study for simplicity. The notation $\text{SNN}(N_{\text{nodes}}, N_{\text{layers}})$ is used to denote SNN model with N_{layers} hidden layers and N_{nodes} nodes per layer, where the activation function is always Tanh , and batch normalization is applied for each layer. Therefore, $\text{GMP}(3,9)$ is a feature set constructed using 9 radial Gaussians and all 7 groups of MCSH from order 0 to 3 (resulting in 63 features) and $\text{GMP}(3,9)+\text{SNN}(50,3)$ is a SNN model based on the $\text{GMP}(3,9)$ feature set with 3 hidden layers and 50 nodes per layer. We note that models with the same number of radial Gaussian probes are not necessarily equivalent, since Gaussian widths are determined manually at this point. Additional details including the widths used for the radial Gaussians are provided in the supplementary information.

First, we compare the GMP+SNN model to the Behler-Parinello neural network (BPNN) approach that is one of the most computationally efficient featurization schemes [39] and ubiquitous in materials science and chemistry due to its simplicity, generality, and efficiency [24, 54, 55, 56, 57, 58, 53]. We utilize an established molecular dynamics trajectory of the 3-element (C, H, O) aspirin molecule at the DFT/PBE+vdW-TS level of theory for this comparison [59]. For BPNN featurization, we use 12 variants inspired by examples in literature [60]. All neural networks consist of 3 hidden layers with 50 nodes each, and the BPNN approach uses separate neural networks for each element (resulting in 3 times the number of parameters). We use a training set size of 40K images for all models, and the test and validation sets each contain 10K images. We ensure robustness by using ten randomly selected train/test/validation sets. The performance is measured by the mean absolute error (MAE) for the predicted energies of the test set images.

Fig 2a shows the GMP+SNN model error as a function of the multipole expansion order and the number of Gaussian radial probes. The results reveal that the GMP+SNN accuracy increases systematically with the multipole expansion order and the number of radial probes, providing a clear strategy for identifying an appropriate feature set for a given problem. In Fig 2b, we compare the accuracy of the GMP+SNN model to the BPNN model as a function of the number of features. All Pareto-optimal models utilize the GMP+SNN approach, despite the fact that the BPNN models have three times more fitted parameters due to element-specific neural networks. It is also clear that the accuracy of BPNN models does not systematically improve with the number of features. Finally, we compare the wall time needed to compute a single image for each model. To ensure a fair comparison, we use the same CPU for each test, both approaches use the same C source code and loop structure, and the time is an average over the 10K predicted validation images (see supplementary information). Fig. 2c) shows that the GMP+SNN framework is always faster than the BPNN approach at a fixed accuracy level, or is always more accurate for a fixed computation time. This example demonstrates that the GMP+SNN approach is capable of achieving a lower error than the BPNN approach with fewer parameters and less time.

The performance of the GMP+SNN can be further improved with force regularization, with a loss function defined as

$$\Gamma = \frac{1}{N_{\text{image}}} \sum_{i=1}^{N_{\text{image}}} \left\{ \|E_{\text{NN}}^i - E_{\text{ref}}^i\| + \frac{\beta}{3N_{\text{atom}}^i} \sum_{j=1}^{3N_{\text{atom}}^i} \|F_{j,\text{NN}}^i - F_{j,\text{ref}}^i\| \right\}, \quad (5)$$

where N_{image} is the number of training images, N_{atom}^i is the number of atoms in image i , E_{NN} and E_{ref} are the predicted and reference energy of a image, F_{NN} and F_{ref} are the predicted and reference forces of an atom along each Cartesian axis, and β is the force regularization coefficient. Figure 3 shows the results of GMP+SNN with 2 feature sets, $\text{GMP}(2,4)+\text{SNN}(50,3)$ and $\text{GMP}(3,12)+\text{SNN}(50,3)$, as a function of the force regularization coefficient, β . For simplicity, this test was done with 1K training images and 1K test images. The results confirm that energy training benefits from force regularization. In this case, the energy error is reduced by a factor of 2 at the optimal regularization coefficient. Both energy and force prediction accuracy show the same trend and the same optimal force coefficient, indicating that the model avoids the apparent trade-off between energy and force accuracy that has been observed for some GCN model architectures [46].

Next, we show that the GMP+SNN technique can scale to systems with more elements by training it on the atomization energy of the established QM9 benchmark dataset [61]. This dataset consists of 134K chemical species with up to 9 heavy atoms and five elements (C, H, O, N, F) optimized at the B3LYP level of theory. In this example, 4 GMP descriptor sets and corresponding SNNs of different sizes are trained and tested using energies of each system. The learning curves showing the out-of-sample test error [62] as a function of training set size are shown in Fig. 4a. The

Figure 2: Training results of GMP+SNN models compared to Behler-Parrinello + HDNN models. The training set consists of 40K images, validation set consists of 10K images and test set consists of 10K images, all randomly drawn from the MD17 aspirin simulation trajectory. Each setting was tested 10 times to ensure robustness. a) Convergence of test MAE as rotational order and number of radial Gaussians are increased with curves for 1 (orange), 2 (green), 3 (red), 4 (purple), 5 (brown) and 6 (pink) radial Gaussians. b) Comparison of test MAE for GMP+SNN (blue) and BPNN (orange) models as a function of the number of features. c) Comparison of test MAE for GMP+SNN and BPNN models as a function of the average prediction time for a single image.

errors of the tested GMP+SNN models are slightly higher than the best state-of-the-art models based on kernel ridge regression or GCNs (about 6 meV with 100K training data) [62, 30, 31, 27, 25, 28, 35, 36], but the GMP+SNN models utilize fewer adjustable parameters and are more scalable than non-parametric models like Gaussian process regression.

Fig. 4b shows the results of transfer learning between different elemental species in the QM9 data set. We trained a model on the set of all molecules that do not contain fluorine (128,908 molecules), and tested the same model on fluorine-containing molecules (1,923 molecules). The model error increases by around one order of magnitude when the predictions include elements not present in the training set. However, the test set MAE is 0.25 eV, which is competitive with the accuracy of generalized gradient approximation of density functional theory (GGA DFT) [63]. This reveals that the model is reasonably accurate even for elements outside the training set. The error distribution for transfer learning is bimodal, which is primarily related to the number of F atoms in the system. If the error is normalized by the number of F atoms in the system the MAE decreases to 0.137 eV/F atom, and the distribution becomes more symmetric (see SI). In addition, we note that the error distribution for F atoms is relatively sensitive to the initial weights of the neural network. The MAE varies from 0.2 - 0.4 eV depending on initial weights, and the distribution is not always bimodal (see SI). This indicates that the model is extrapolating and therefore more sensitive to the initial conditions and

Figure 3: Test MAE of energies and forces as a function of the force regularization coefficient on 1K randomly selected images from the aspirin trajectory. Two feature sets are compared here, both with SNN(50,3) architecture: GMP(2,4) (red, 16 features) and GMP(3,12) blue, 84 features).

Figure 4: Learning curve and elemental transferability of GMP+SNN on the QM9 dataset. a) Learning curves of the GMP+SNN models with various feature sets and neural net sizes (blue = GMP(1,8)+SNN(15,5), green = GMP(2,8)+SNN(25,5), yellow = GMP(2,13)+SNN(60,5), red = GMP(2,37)+SNN(100,5)). The dashed line indicates chemical accuracy (43 meV). b) Test MAE distribution for model trained exclusively on molecules without F and tested on all F-containing molecules (blue) and test MAE distribution for model trained and tested with randomly selected molecules (red).

training procedure. However, the resulting models never fail catastrophically since the errors are always competitive with GGA DFT, indicating the general robustness of the transfer learning between elements.

Finally, we test the performance of the GMP+SNN approach on the Open Catalysis Project (OC20) S2EF dataset. This dataset consists of >100 M geometries and corresponding energies of >100 K adsorbate-catalyst pairs containing a total of 56 elements across 82 adsorbates and up to >100 K different catalyst compositions for each adsorbate. The dataset also provides an independent set of 1M test systems [46]. The energies correspond to the GGA DFT level of theory with the RBPE functional [64], with mixed boundary conditions. The size, number of elements, and mix of solid-state and molecular systems make this one of the most challenging benchmark datasets available. To date, the only models capable of training and prediction for this dataset utilize elaborate GCN models [46]. Fig. 5a shows the learning curves for three GMP+SNN models of different sizes tested on the provided in-domain (ID) validation set. The small model uses the GMP(3,13)+SNN(80,6) architecture (91 features, 34K parameters), the medium model uses the GMP(4,19)+SNN(150,7) architecture (209 features, 146K parameters), and the large model

uses the GMP(5,37)+SNN(250,8) architecture (592 features, 528K parameters). Details of all models are provided in the supplementary information. The full OC20 training set of 100M energies would require months to train with current computing resources, so we restrict the analysis to energy training on data sets with fewer than 5M training images.

The results of training and testing on the OC20 set are shown in Fig. 5a. The in-domain test error reaches a minimum of 0.56 eV with 5M training images with the GMP(5,37)+SNN(250,8) model, an error lower than all GCN models with 5M training points, and comparable to the performance of the CGCNN (0.527 eV) and DimeNet++ (0.486 eV) GCN models trained on all >100 M training points [46]. Moreover, the GMP+SNN models require far fewer parameters than the GCN models, with the largest GMP+SNN model having ~ 500 K parameters, compared to 1.8M (DimeNet++), 3.6M (CGCNN), and 7.4M (Schnet) parameters for the GCN models[46]. The error also decreases as the number of data points, the number of GMP features, or the neural net size increase. This suggests that further improvement is possible, although significant computational resources will be required to optimize and evaluate GMP+SNN models for the OC20 dataset.

Finally, the scaling of computational time versus the number of atoms in the system is plotted in Fig. 5b. The systems of different sizes are generated by making supercells that are periodic replications of an original unit cell of 35 atoms. The scaling is strictly linear due to the local nature of the features, which presents an advantage over most graph-based models require the entire structure as input. The linear scaling of the GMP+SNN model, along with a structure that is straightforward to parallelize, makes it ideal for the simulation of large systems that may not be feasible with graph-based models.

Figure 5: Results for the OC20 dataset with GMP(3,13)+SNN(80,6) (blue), GMP(4,19)+SNN(150,7) (purple), and GMP(5,37)+SNN(250,8) (red) models. a) Learning curves tested on the provided in-domain (ID) validation set with the DimeNet++ (gray dashed) results for comparison [46] b) Log-log plot of time required to infer one images v.s. the number of atoms present. Systems of different sizes are made by periodically repeating one unit cell of C_2O adsorbed on Sn_3Ca (4 elements, 35 atoms).

These examples demonstrate that the GMP featurization scheme is an efficient and universal approach to fingerprinting atomistic systems with an arbitrary number of elements or atom types. The GMP features are more computationally efficient than the deep-learning GCN models that are commonly used for many-element systems, and are even faster than the widely used Behler-Parinello symmetry functions, resulting in machine-learned force fields that can be scaled to large many-element systems and long time scales. The GMP feature vectors utilize an implicit description of the system’s electron density, making the number of features independent of the number of elements, facilitating the inclusion of general concepts from electronic structure theory, and enabling extrapolation to new element types. Moreover, the GMP features utilize physically-meaningful concepts, leading to more interpretable models that can be systematically improved and providing a foundation for hybrid models that incorporate more physics. For example, the self-consistent electron density from a simplified Hamiltonian could be used as input to the GMP model, or the limiting behavior derived from electrostatics could be used to constrain the regression model.

The examples presented here use the SNN regression model in conjunction with the GMP features based on static valence density to obtain accuracy comparable to state-of-the-art models on the MD17, QM9, and OC20 datasets. The results confirm that the GMP+SNN approach can reach competitive accuracy with GCN based models despite having far fewer adjustable parameters and a simpler structure. There are numerous opportunities to revise and optimize the details of the GMP+SNN approach presented here, including modifying the valence density representation, optimizing feature selection and systematic optimization of the SNN architecture. However, the encouraging initial results across a

wide range of application domains suggest that the GMP+SNN approach is a promising universal route to constructing efficient and general machine-learning models for atomistic systems.

Acknowledgments

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences Computational Chemical Sciences program under Award Numbers DE-SC0019441 and DE-SC0019410. Computational effort was supplied partially by the National Science Foundation under Grant No. MRI-1828187. We acknowledge helpful discussions with Andrew Peterson, Zachary Ulissi, and Muhammed Shuaibi.

References

- [1] Kieron Burke. Perspective on density functional theory. *The Journal of Chemical Physics*, 136(15):150901, 2012.
- [2] Axel D. Becke. Perspective: Fifty years of density-functional theory in chemical physics. *The Journal of Chemical Physics*, 140(18):18A301, 2014.
- [3] Narbe Mardirossian and Martin Head-Gordon. Thirty years of density functional theory in computational chemistry: an overview and extensive assessment of 200 density functionals. *Molecular physics*, 115(19):2315–2372, 2017.
- [4] Adam Hospital, Josep Ramon Goñi, Modesto Orozco, and Josep L Gelpí. Molecular dynamics simulations: advances and applications. *Advances and applications in bioinformatics and chemistry*, 8:37–47, 2015.
- [5] Marco De Vivo, Matteo Masetti, Giovanni Bottegoni, and Andrea Cavalli. Role of molecular dynamics and related methods in drug discovery. *Journal of medicinal chemistry*, 59(9):4035–4061, 2016.
- [6] J. Andrew McCammon and Martin Karplus. Molecular dynamics simulations of biomolecules. *Nature structural biology*, 9(9):646–652, 2002.
- [7] Judith A Harrison, J. David Schall, Sabina Maskey, Paul T Mikulski, M. Todd Knippenberg, and Brian H Morrow. Review of force fields and intermolecular potentials used in atomistic computational materials research. *Applied Physics Reviews*, 5(3):31104, 2018.
- [8] Artem R Oganov, Chris J Pickard, Qiang Zhu, and Richard J Needs. Structure prediction drives materials discovery. *Nature reviews. Materials*, 4(5):331–348, 2019.
- [9] Xinguo Ren, Patrick Rinke, Christian Joas, and Matthias Scheffler. Random-phase approximation and its applications in computational chemistry and materials science. *Journal of materials science*, 47(21):7447–7471, 2012.
- [10] Edoardo Mosconi, Jon M Azpiroz, and Filippo De Angelis. Ab initio molecular dynamics simulations of methylammonium lead iodide perovskite degradation by water. *Chemistry of materials*, 27(13):4885–4892, 2015.
- [11] Dominik Marx and Jürg Hutter. *Ab Initio Molecular Dynamics: Basic Theory and Advanced Methods*. Cambridge University Press, 2009.
- [12] Radu Iftimie, Peter Minari, Mark E. Tuckerman, and Bruce J. Berne. Ab initio molecular dynamics: Concepts, recent developments, and future trends. *Proceedings of the National Academy of Sciences - PNAS*, 102(19):6654–6659, 2005.
- [13] B. R Brooks, C. L Brooks, A. D Mackerell, L Nilsson, R. J Petrella, B Roux, Y Won, G Archontis, C Bartels, S Boresch, A Caffisch, L Caves, Q Cui, A. R Dinner, M Feig, S Fischer, J Gao, M Hodoscek, W Im, K Kuczera, T Lazaridis, J Ma, V Ovchinnikov, E Paci, R. W Pastor, C. B Post, J. Z Pu, M Schaefer, B Tidor, R. M Venable, H. L Woodcock, X Wu, W Yang, D. M York, and M Karplus. Charmm: The biomolecular simulation program. *Journal of computational chemistry*, 30(10):1545–1614, 2009.
- [14] David A Case, Thomas E Cheatham, Tom Darden, Holger Gohlke, Ray Luo, Kenneth M Merz, Alexey Onufriev, Carlos Simmerling, Bing Wang, and Robert J Woods. The amber biomolecular simulation programs. *Journal of computational chemistry*, 26(16):1668–1688, 2005.
- [15] Adri C. T van Duin, Siddharth Dasgupta, Francois Lorant, and William A Goddard. Reaxff: A reactive force field for hydrocarbons. *The journal of physical chemistry. A, Molecules, spectroscopy, kinetics, environment, and general theory*, 105(41):9396–9409, 2001.
- [16] William L Jorgensen and Julian Tirado-Rives. The opls [optimized potentials for liquid simulations] potential functions for proteins, energy minimizations for crystals of cyclic peptides and crambin. *Journal of the American Chemical Society*, 110(6):1657–1666, 1988.

- [17] Simon R. Phillpot, Andrew C. Antony, Linyuan Shi, Michele L. Fullarton, Tao Liang, Susan B. Sinnott, Yongfeng Zhang, and S. Bulent Biner. Charge optimized many body (comb) potentials for simulation of nuclear fuel and clad. *Computational Materials Science*, 148:231 – 241, 2018.
- [18] Thomas P Senftle, Sungwook Hong, Md Mahbulul Islam, Sudhir B Kylasa, Yuanxia Zheng, Yun Kyung Shin, Chad Junkermeier, Roman Engel-Herbert, Michael J Janik, Hasan Metin Aktulga, Toon Verstraelen, Ananth Grama, and Adri C T van Duin. The reaxff reactive force-field: development, applications and future directions. *npj computational materials*, 2(1):15011, 2016.
- [19] Alexander D MacKerell. Chapter 7 empirical force fields for proteins: Current status and future directions. *Annual Reports in Computational Chemistry*, 1:91–102, 2005.
- [20] Judith A Harrison, J. David Schall, Sabina Maskey, Paul T Mikulski, M. Todd Knippenberg, and Brian H Morrow. Review of force fields and intermolecular potentials used in atomistic computational materials research. *Applied Physics Reviews*, 5(3):31104, 2018.
- [21] Chris M Handley and Paul L. A Popelier. Potential energy surfaces fitted by artificial neural networks. *The journal of physical chemistry. A, Molecules, spectroscopy, kinetics, environment, and general theory*, 114(10):3371–3383, 2010.
- [22] Jörg Behler. Perspective: Machine learning potentials for atomistic simulations. *The Journal of chemical physics*, 145(17):170901–170901, 2016.
- [23] Rampi Ramprasad, Rohit Batra, Ghanshyam Pilania, Arun Mannodi-Kanakkithodi, and Chiho Kim. Machine learning in materials informatics: recent applications and prospects. *npj computational materials*, 3(1):1–13, 2017.
- [24] Jörg Behler and Michele Parrinello. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.*, 98:146401, 4 2007.
- [25] Kristof Schütt, Pieter-Jan Kindermans, Huziel Enoc Saucedo Felix, Stefan Chmiela, Alexandre Tkatchenko, and Klaus-Robert Müller. Schnet: A continuous-filter convolutional neural network for modeling quantum interactions. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30, pages 991–1001. Curran Associates, Inc., 2017.
- [26] 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):145301–145301, 2018.
- [27] Johannes Klicpera, Janek Groß, and Stephan Günnemann. Directional message passing for molecular graphs. In *International Conference on Learning Representations*, 2020.
- [28] Oliver T Unke and Markus Meuwly. Physnet: A neural network for predicting energies, forces, dipole moments, and partial charges. *Journal of chemical theory and computation*, 15(6):3678–3693, 2019.
- [29] Nicholas Lubbers, Justin S Smith, and Kipton Barros. Hierarchical modeling of molecular energies using a deep neural network. *The Journal of chemical physics*, 148(24):241715–241715, 2018.
- [30] Shuo Zhang, Yang Liu, and Lei Xie. Molecular mechanics-driven graph neural network with multiplex graph for molecular structures. 2020.
- [31] Zeren Shui and George Karypis. Heterogeneous molecular graph neural networks for predicting molecule properties. 2020.
- [32] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. 2017.
- [33] Mingjie Liu and John R. Kitchin. Singlenn: Modified behler–parrinello neural network with shared weights for atomistic simulations with transferability. *The Journal of Physical Chemistry C*, 124(32):17811–17818, 2020.
- [34] J. S. Smith, O. Isayev, and A. E. Roitberg. Ani-1: an extensible neural network potential with dft accuracy at force field computational cost. *Chem. Sci.*, 8:3192–3203, 2017.
- [35] Felix A Faber, Anders S Christensen, Bing Huang, and O. Anatole von Lilienfeld. Alchemical and structural distribution based representation for universal quantum machine learning. *The Journal of chemical physics*, 148(24):241717–241717, 2018.
- [36] Albert P Bartók, Sandip De, Carl Poelking, Noam Bernstein, James R Kermode, Gábor Csányi, and Michele Ceriotti. Machine learning unifies the modeling of materials and molecules. *Science advances*, 3(12):e1701816–e1701816, 2017.
- [37] Miguel A. Caro. Optimizing many-body atomic descriptors for enhanced computational performance of machine learning based interatomic potentials. *Phys. Rev. B*, 100:024112, Jul 2019.

- [38] Michael Eickenberg, Georgios Exarchakis, Matthew Hirn, and Stephane Mallat. Solid harmonic wavelet scattering: Predicting quantum molecular energy from invariant descriptors of 3d electronic densities. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [39] Marcel F. Langer, Alex Goeßmann, and Matthias Rupp. Representations of molecules and materials for interpolation of quantum-mechanical simulations via machine learning. *arXiv*, 2020.
- [40] Felix Musil, Andrea Grisafi, Albert P. Bartók, Christoph Ortner, Gábor Csányi, and Michele Ceriotti. Physics-inspired structural representations for molecules and materials. *Chemical Reviews*, 121(16):9759–9815, jul 2021.
- [41] Matthias Rupp, Alexandre Tkatchenko, Klaus-Robert Müller, and O Anatole von Lilienfeld. Fast and accurate modeling of molecular atomization energies with machine learning. *Physical review letters*, 108(5):058301–058301, 2012.
- [42] Venkatesh Botu and Rampi Ramprasad. Adaptive machine learning framework to accelerate ab initio molecular dynamics. *International journal of quantum chemistry*, 115(16):1074–1083, 2015.
- [43] Nongnuch Artrith, Alexander Urban, and Gerbrand Ceder. Efficient and accurate machine-learning interpolation of atomic energies in compositions with many species. *Phys. Rev. B*, 96:014112, Jul 2017.
- [44] J Behler. Representing potential energy surfaces by high-dimensional neural network potentials. *Journal of Physics: Condensed Matter*, 26(18):183001, apr 2014.
- [45] Linfeng Zhang, Jiequn Han, Han Wang, Roberto Car, and Weinan E. Deep potential molecular dynamics: A scalable model with the accuracy of quantum mechanics. *Physical Review Letters*, 120(14), apr 2018.
- [46] 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, may 2021.
- [47] A. R Edmonds. *Angular momentum in quantum mechanics*. Investigations in physics ; 4. 1957.
- [48] William J Thompson and LeRoy F Cook. Angular momentum: An illustrated guide to rotational symmetries for physical systems. *American Journal of Physics*, 63(7):670–671, 1995.
- [49] D. R. Hamann. Optimized norm-conserving vanderbilt pseudopotentials. *Phys. Rev. B*, 88:085117, Aug 2013.
- [50] Xiangyun Lei and Andrew J. Medford. Design and analysis of machine learning exchange-correlation functionals via rotationally invariant convolutional descriptors. *Phys. Rev. Materials*, 3:063801, 6 2019.
- [51] Jon Applequist. Maxwell–cartesian spherical harmonics in multipole potentials and atomic orbitals. *Theoretical Chemistry Accounts*, 107(2):103–115, 2002.
- [52] Amptorch. <https://github.com/ulissigroup/amptorch>, 2020.
- [53] Muhammed Shuaibi, Saurabh Sivakumar, Rui Qi Chen, and Zachary W Ulissi. Enabling robust offline active learning for machine learning potentials using simple physics-based priors. *Machine Learning: Science and Technology*, 2(2), 2020.
- [54] Kyuhyun Lee, Dongsun Yoo, Wonseok Jeong, and Seungwu Han. Simple-nn: An efficient package for training and executing neural-network interatomic potentials. *Computer Physics Communications*, 242:95 – 103, 2019.
- [55] Alireza Khorshidi and Andrew A. Peterson. Amp: A modular approach to machine learning in atomistic simulations. *Computer Physics Communications*, 207:310 – 324, 2016.
- [56] Nongnuch Artrith, Tobias Morawietz, and Jörg Behler. High-dimensional neural-network potentials for multicomponent systems: Applications to zinc oxide. *Phys. Rev. B*, 83:153101, Apr 2011.
- [57] Nongnuch Artrith and Jörg Behler. High-dimensional neural network potentials for metal surfaces: A prototype study for copper. *Phys. Rev. B*, 85:045439, Jan 2012.
- [58] Nongnuch Artrith, Björn Hiller, and Jörg Behler. Neural network potentials for metals and oxides – first applications to copper clusters at zinc oxide. *physica status solidi (b)*, 250(6):1191–1203, 2013.
- [59] 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):3887, 2018.
- [60] Christoph Schran, Jörg Behler, and Dominik Marx. Automated fitting of neural network potentials at coupled cluster accuracy: Protonated water clusters as testing ground. *Journal of chemical theory and computation*, 16(1):88–99, 2020.

- [61] 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):140022–140022, 2014.
- [62] O. Anatole von Lilienfeld, Klaus-Robert Müller, and Alexandre Tkatchenko. Exploring chemical compound space with quantum-based machine learning. 2019.
- [63] Kevin E Riley, Bryan T Op’t Holt, and Kenneth M Merz. Critical assessment of the performance of density functional methods for several atomic and molecular properties. *Journal of chemical theory and computation*, 3(2):407–433, 2007.
- [64] B. Hammer, L. B. Hansen, and J. K. Nørskov. Improved adsorption energetics within density-functional theory using revised perdew-burke-ernzerhof functionals. *Phys. Rev. B*, 59:7413–7421, Mar 1999.

Supplementary Information for “A Universal Framework for Featurization of Atomistic Systems”

0.1 Code and Saved Models

Please checkout and install the “MCSH_paper1” branch of *AMPTorch* to try any of the saved models and test scripts in this study. The code can be found here: https://github.com/ulissigroup/amptorch/tree/MCSH_paper1. The test scripts can be found here: https://github.com/ray38/GMP_AmpTorch_Tests. Tutorials for regenerating all the test models are all included in the repo.

0.2 MD17 Aspirin Examples

0.2.1 Standard Training Procedure

Cutoffs for both the BP scheme and the GMP scheme are set to 10 Å. The neural networks are all trained using the same procedure: $lr = 1e^{-3}$ for 6000 epochs. For the BP vs. GMP comparison example on aspirin MD data, the batch size is chosen to be 256 images. For the force training example, the batch size is chosen to be 32 images.

0.2.2 Behler-Parrinello + HDNN Comparison Test Setups

12 sets of Behler-Parrinello feature are selected, as listed below

Set	G2		G4			$N_{feature}^a$
	η	R_s	η	ζ	γ	
1	[0.05, 0.0965, 0.1864, 0.3598, 0.6947, 1.3413, 2.5897, 5.0]	[0, 1.5]	[0.001, 0.01, 0.03]	[1.0, 2.0, 4.0]	[1.0, -1.0]	156
2	[0.05, 0.0965, 0.1864, 0.3598, 0.6947, 1.3413, 2.5897, 5.0]	[0, 1.5]	[0.01, 0.03]	[1.0, 4.0]	[1.0, -1.0]	96
3	[0.05, 0.0965, 0.1864, 0.3598, 0.6947, 1.3413, 2.5897, 5.0]	[0]	[0.01]	[1.0, 4.0]	[1.0, -1.0]	48
4	[0.05, 0.0965, 0.1864, 0.3598, 0.6947, 1.3413, 2.5897, 5.0]	[0]	[0.001, 0.01, 0.03]	[1.0, 2.0, 4.0]	[1.0, -1.0]	132
5	[0.05, 0.0965, 0.1864, 0.3598, 0.6947, 1.3413, 2.5897, 5.0]	[0]	[0.01, 0.03]	[1.0, 4.0]	[1.0, -1.0]	72
6	[0.05, 0.0965, 0.1864, 0.3598, 0.6947, 1.3413, 2.5897, 5.0]	[0, 1.5]	[0.01]	[1.0, 4.0]	[1.0, -1.0]	72
7	[0.05, 0.2324, 1.0772, 5.]	[0, 1.5]	[0.001, 0.01, 0.03]	[1.0, 2.0, 4.0]	[1.0, -1.0]	132
8	[0.05, 0.2324, 1.0772, 5.]	[0, 1.5]	[0.01, 0.03]	[1.0, 4.0]	[1.0, -1.0]	72
9	[0.05, 0.2324, 1.0772, 5.]	[0]	[0.01]	[1.0, 4.0]	[1.0, -1.0]	32
10	[0.05, 0.2324, 1.0772, 5.]	[0]	[0.001, 0.01, 0.03]	[1.0, 2.0, 4.0]	[1.0, -1.0]	120
11	[0.05, 0.2324, 1.0772, 5.]	[0]	[0.01, 0.03]	[1.0, 4.0]	[1.0, -1.0]	60
12	[0.05, 0.2324, 1.0772, 5.]	[0, 1.5]	[0.01]	[1.0, 4.0]	[1.0, -1.0]	48

Table 2: List of tested Behler-Parrinello feature sets. η and R_s for G2 functions are used combinatorially, same as η , ζ and γ for G4 functions. Moreover, there are 3 types of elements (C, H, O) for this dataset. Therefore, feature set 1 has $3(elements) \times 8(\eta) \times 2(R_s) + 6(possible\ element\ pairs) \times 3(\eta) \times 3(\zeta)s \times 2(\gamma) = 156$ features per atom.

^aNumber of features per atom.

The list of test results with BP + HDNN models are shown below in Table 3

Set	$N_{feature}^a$	MAE train (meV)	MAE test (meV)	Time (ms/image)
1	156	73.0 ± 3.0	98.8 ± 2.7	22.1 ± 0.5
2	96	74.2 ± 1.4	98.8 ± 1.7	13.3 ± 0.4
3	48	71.8 ± 1.9	97.3 ± 1.8	8.2 ± 0.1
4	132	73.5 ± 1.6	100.3 ± 2.5	21.2 ± 0.3
5	72	73.2 ± 1.2	99.6 ± 1.5	12.1 ± 0.1
6	72	74.4 ± 1.6	97.4 ± 2.2	9.1 ± 0.4
7	132	73.3 ± 1.4	99.9 ± 1.5	20.9 ± 0.1
8	72	71.6 ± 1.5	97.8 ± 2.4	12.1 ± 0.1
9	32	70.5 ± 2.2	97.3 ± 1.9	8.1 ± 0.2
10	120	73.5 ± 1.5	101.2 ± 1.4	20.4 ± 0.0
11	60	71.7 ± 2.1	99.3 ± 1.6	11.8 ± 0.1
12	48	72.4 ± 1.7	96.9 ± 1.3	8.2 ± 0.1

Table 3: Performance test results of the tested GMP + HDNN setups. The values are the average values of the 10 trials, and the uncertainties are estimated by their standard deviation. ^aNumber of features per atom.

0.2.3 GMP+SNN Comparison Test Setups

The probe of GMP has two parts: groups of MCSHs for probing angular features and radial gaussian for probing radial features. In this work, all groups of MCSH up to the indicated order are included, and the number of possible groups for each order are listed below in Table 4:

Order	Number of Possible Groups	Total Number of Groups up to This Order
0	1	1
1	1	2
2	2	4
3	3	7
4	4	11
5	5	16
6	7	23
7	8	31
8	10	41

Table 4: List of possible groups of MCSHs for each order. Note that the number of possible groups

We combine the radial probes with the lists of radial Gaussians combinatorially to obtain the full list of probes/features. The list of widths (standard deviations) of the Gaussians used in this test are manually picked and listed below in Table 5:

Number of Gaussians	List of Sigmas
1	[0.25]
2	[0.25, 2.0]
3	[0.25, 1.0, 2.0]
4	[0.25, 0.75, 1.5, 2.0]
5	[0.25, 0.5, 1.0, 1.5, 2.0]
6	[0.25, 0.5, 0.75, 1.0, 1.5, 2.0]

Table 5: List of Sigmas used in Test 1.

Therefore, when there are 5 Gaussians with MCSH up to order 6, there are $23 \times 5 = 115$ descriptors per atom. The complete list of test results with GMP+SNN is give in Table 6.

Num. Gaussians ^a	MCSH order ^b	$N_{feature}$ ^c	MAE train (meV)	MAE test (meV)	Time (ms/image)
1	0	1	180.1 $\pm$ 0.5	180.7 $\pm$ 0.9	2.7 $\pm$ 0.2
	1	2	146.5 $\pm$ 0.9	146.9 $\pm$ 1.1	2.8 $\pm$ 0.2
	2	4	133.9 $\pm$ 0.9	134.6 $\pm$ 1.1	3.0 $\pm$ 0.2
	3	7	116.8 $\pm$ 1.3	119.5 $\pm$ 1.6	3.1 $\pm$ 0.2
	4	11	114.7 $\pm$ 1.6	117.2 $\pm$ 1.3	3.4 $\pm$ 0.2
	5	16	104.1 $\pm$ 2.1	107.8 $\pm$ 1.7	3.8 $\pm$ 0.2
	6	23	103.2 $\pm$ 1.1	107.1 $\pm$ 1.3	4.1 $\pm$ 0.1
	7	31	95.2 $\pm$ 1.1	99.4 $\pm$ 1.6	4.8 $\pm$ 0.2
2	0	2	162.4 $\pm$ 2.1	163.5 $\pm$ 2.1	2.7 $\pm$ 0.1
	1	4	117.5 $\pm$ 1.5	120.5 $\pm$ 1.7	2.9 $\pm$ 0.2
	2	8	100.0 $\pm$ 0.8	105.9 $\pm$ 1.5	3.2 $\pm$ 0.2
	3	14	88.3 $\pm$ 1.9	98.1 $\pm$ 1.4	3.5 $\pm$ 0.2
	4	22	87.1 $\pm$ 1.8	97.7 $\pm$ 2.0	3.9 $\pm$ 0.1
	5	32	84.3 $\pm$ 0.9	97.0 $\pm$ 1.2	4.7 $\pm$ 0.2
	6	46	80.7 $\pm$ 1.0	94.9 $\pm$ 1.2	5.9 $\pm$ 0.4
	7	62	74.8 $\pm$ 1.2	88.1 $\pm$ 1.2	7.1 $\pm$ 0.3
3	0	3	142.7 $\pm$ 2.1	144.8 $\pm$ 1.6	2.7 $\pm$ 0.2
	1	6	89.2 $\pm$ 1.0	92.5 $\pm$ 0.9	3.0 $\pm$ 0.4
	2	12	78.2 $\pm$ 0.8	84.1 $\pm$ 0.7	3.3 $\pm$ 0.2
	3	21	74.1 $\pm$ 0.9	83.8 $\pm$ 1.4	4.0 $\pm$ 0.3
	4	33	72.3 $\pm$ 0.8	83.6 $\pm$ 0.9	4.9 $\pm$ 0.3
	5	48	67.9 $\pm$ 1.4	81.1 $\pm$ 2.2	6.0 $\pm$ 0.3
	6	69	65.5 $\pm$ 1.6	80.0 $\pm$ 1.7	7.7 $\pm$ 0.2
	7	93	60.4 $\pm$ 0.8	76.3 $\pm$ 1.7	9.8 $\pm$ 0.4
4	0	4	120.3 $\pm$ 1.7	121.6 $\pm$ 1.6	2.9 $\pm$ 0.2
	1	8	77.2 $\pm$ 1.3	79.8 $\pm$ 1.0	3.0 $\pm$ 0.2
	2	16	69.9 $\pm$ 0.9	75.3 $\pm$ 0.9	3.4 $\pm$ 0.0
	3	28	66.2 $\pm$ 2.2	75.3 $\pm$ 1.8	4.4 $\pm$ 0.2
	4	44	64.5 $\pm$ 0.9	75.8 $\pm$ 1.6	5.6 $\pm$ 0.3
	5	64	61.8 $\pm$ 0.6	75.2 $\pm$ 1.0	7.1 $\pm$ 0.4
	6	92	57.8 $\pm$ 1.0	73.0 $\pm$ 1.2	9.5 $\pm$ 0.3
	7	124	54.5 $\pm$ 0.7	70.8 $\pm$ 1.2	11.8 $\pm$ 0.5
5	0	5	92.3 $\pm$ 1.3	93.3 $\pm$ 1.3	2.9 $\pm$ 0.1
	1	10	62.6 $\pm$ 1.1	65.1 $\pm$ 1.0	3.1 $\pm$ 0.2
	2	20	62.0 $\pm$ 1.9	66.3 $\pm$ 1.8	3.6 $\pm$ 0.0
	3	35	59.4 $\pm$ 1.4	67.1 $\pm$ 1.6	4.7 $\pm$ 0.1
	4	55	56.1 $\pm$ 1.1	65.3 $\pm$ 1.1	6.4 $\pm$ 0.3
	5	80	50.9 $\pm$ 0.5	61.5 $\pm$ 0.9	8.3 $\pm$ 0.3
	6	115	47.3 $\pm$ 0.8	60.1 $\pm$ 1.1	10.8 $\pm$ 0.3
	7	155	43.8 $\pm$ 0.8	57.6 $\pm$ 1.4	14.0 $\pm$ 0.4
6	0	6	85.3 $\pm$ 1.8	86.5 $\pm$ 1.8	2.9 $\pm$ 0.2
	1	12	60.6 $\pm$ 1.5	62.9 $\pm$ 1.3	3.5 $\pm$ 0.2
	2	24	58.3 $\pm$ 1.2	62.5 $\pm$ 1.4	4.2 $\pm$ 0.3
	3	42	54.6 $\pm$ 1.0	61.5 $\pm$ 0.7	5.3 $\pm$ 0.3
	4	66	51.7 $\pm$ 0.8	60.7 $\pm$ 1.4	7.0 $\pm$ 0.4
	5	96	47.3 $\pm$ 0.9	58.7 $\pm$ 1.2	9.1 $\pm$ 0.4
	6	138	44.0 $\pm$ 0.7	56.4 $\pm$ 0.9	12.2 $\pm$ 0.6
	7	186	41.3 $\pm$ 0.9	55.7 $\pm$ 0.9	15.8 $\pm$ 0.6

Table 6: Performance test results of the full sets of tested GMP + SNN setups. The values are the average values of the 10 trials, and the uncertainties are estimated by their standard deviation. ^aNumber of possible Gaussian functions used to construct the descriptor probes. ^bThe highest MCSH order used to construct the probes. For example, when highest order is 2, that means all groups from MCSH of order 0, 1 and 2 are used to construct the probes. ^cNumber of features per atom.

Model	Sigmas	$N_{feature}$
$GMP(2, 4) + SNN(50, 3)$	[0.25, 0.75, 1.5, 2.0]	16
$GMP(3, 12) + SNN(50, 3)$	[0.125, 0.25, 0.375, 0.5, 0.625, 0.75, 0.875, 1.0, 1.25, 1.5, 1.75, 2.0]	84

Table 7: Setups for the GMP+SNN models used in the force training example

0.2.4 Force Training Example

The sigmas of the radial probe Gaussians are listed in Table 7

0.3 QM9 Example

0.3.1 Per-element Bias

A per-element bias is added to the SNN model to improve performance. Conceptually, this is equivalent to fitting to formation energies rather than absolute energies. The total energy of an atom is the model predicted energy plus the per-element bias of the specific atom type. To determine the bias, a linear model is applied. The number of atoms for each element types are counted for all the images in the training set, and they are the independent variable. The corresponding energies for each system is the dependent variable. For example, the per-element bias of the trials with 100K training images is shown below in Table 8:

Atom Type	Per-element Bias (meV)
H	-2795.2721
C	-6217.7719
N	-4552.4431
O	-4432.6761
F	-4075.4931

Table 8: Per-element Bias of each atom type found by the linear model, for the 100K training set

0.3.2 Training Procedure

With the per-element bias determined, GMP+SNN models are fitted to the atomization energy minus the per-element biases. The model setups are given in Table 9. The cutoff distance is always 15 Å, so that the largest radial probe takes a negligible value of 8×10^{-4} at the cutoff. The models are trained for 12,000 epochs with learning rate decrease by factor of 2 every 2,000 epochs, from 1^{-2} to 3^{-4} . The batch size is set to be 32 images.

0.3.3 Transfer Learning to New Element

For this example, we used the same procedure as above, with the caveat that the per-element biases are not fitted using a linear model, but directly pulled from the 100k molecule trial. For more detail please refer to the test scripts.

Shown in Figure 6 is the error distribution of the system containing F atoms in the basis of error per F atom. Shown in Figure 7 are the error distributions from different trials with different random seeds.

Model	Sigmas	$N_{feature}$	$N_{parameters}$
$GMP(1, 8) + SNN(15, 5)$	[0.02,0.2,0.4,0.69,1.1,1.66,2.66,4.4]	16	1111
$GMP(2, 8) + SNN(25, 5)$	[0.02,0.2,0.4,0.69,1.1,1.66,2.66,4.4]	84	3001
$GMP(2, 13) + SNN(60, 5)$	[0.02,0.12,0.24,0.36,0.5,0.69,0.92,1.2,1.52,2.0,2.66,3.5,5.0]	52	14701
$GMP(2, 37) + SNN(100, 5)$	[0.02,0.05,0.08,0.12,0.16,0.2,0.24,0.28,0.32,0.36,0.4,0.45,0.5,0.56,0.62,0.69,0.76,0.84,0.92,1.01,1.1,1.2,1.3,1.4,1.52,1.66,1.82,2.0,2.42,2.66,2.92,3.2,3.5,3.9,4.4,5.0]	148	45701

Table 9: Setups for the GMP+SNN models used in the QM9 examples. Cutoff distance is always 15 Å. $N_{parameters}$ is the number of trainable parameters of the neural network model.

Figure 6: Error distribution of F-containing molecules per F atom.

Figure 7: Error distribution of prediction for F-containing molecules with different random seed for the model trained on molecules without F.

Model	Sigmas	$N_{feature}$	$N_{parameters}$
GMP(3,13)+SNN(80,6)	[0.02,0.16,0.32,0.5,0.76,1.1,1.52,2.2,3.2,5.0]	91	34161
GMP(4,19)+SNN(150,7)	[0.02,0.08,0.16,0.24,0.32,0.4,0.5,0.62,0.76,0.92,1.1,1.3,1.52,1.82,2.2,2.66,3.2,3.9,5.0]	209	146701
GMP(5,37)+SNN(250,8)	[0.02,0.05,0.08,0.12,0.16,0.2,0.24,0.28,0.32,0.36,0.4,0.45,0.5,0.56,0.62,0.69,0.76,0.84,0.92,1.01,1.1,1.2,1.3,1.4,1.52,1.66,1.82,2.0,2.42,2.66,2.92,3.2,3.5,3.9,4.4,5.0]	592	528501

Table 10: Setups for the GMP+SNN models used in the OC20 examples. Cutoff distance is always 15 Å. $N_{parameters}$ is the number of trainable parameters (weights) of the neural network model.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [MedfordLeiSI.pdf](#)