

Physical pooling functions in graph neural networks for molecular property prediction

Schweidtmann, Artur M.; Rittig, Jan G.; Weber, Jana M.; Grohe, Martin; Dahmen, Manuel; Leonhard, Kai; Mitsos, Alexander

DOI

[10.1016/j.compchemeng.2023.108202](https://doi.org/10.1016/j.compchemeng.2023.108202)

Publication date

2023

Document Version

Final published version

Published in

Computers and Chemical Engineering

Citation (APA)

Schweidtmann, A. M., Rittig, J. G., Weber, J. M., Grohe, M., Dahmen, M., Leonhard, K., & Mitsos, A. (2023). Physical pooling functions in graph neural networks for molecular property prediction. *Computers and Chemical Engineering*, 172, Article 108202. <https://doi.org/10.1016/j.compchemeng.2023.108202>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



Physical pooling functions in graph neural networks for molecular property prediction

Artur M. Schweidtmann^{a,b}, Jan G. Rittig^a, Jana M. Weber^c, Martin Grohe^d, Manuel Dahmen^e, Kai Leonhard^f, Alexander Mitsos^{g,e,a,*}

^a RWTH Aachen University, Process Systems Engineering (AVT.SVT), Forckenbeckstr. 51, 52074 Aachen, Germany

^b Delft University of Technology, Department of Chemical Engineering, Van der Maasweg 9, Delft 2629 HZ, The Netherlands

^c Delft University of Technology, Delft Bioinformatics Lab, Intelligent Systems, TU Delft, 2628 XE Delft, The Netherlands

^d RWTH Aachen University, Lehrstuhl für Informatik 7, Ahornstr. 55, 52074 Aachen, Germany

^e Forschungszentrum Jülich GmbH, Institute of Energy and Climate Research IEK-10 – Energy Systems Engineering, Wilhelm-Johnen-Str., 52425, Jülich, Germany

^f RWTH Aachen University, Institute of Technical Thermodynamics, Schinkelstr. 8, 52062 Aachen, Germany

^g JARA Center for Simulation and Data Science (CSD), Aachen, Germany

ARTICLE INFO

Keywords:

Graph convolutional neural networks
Pooling function
Physics-informed machine learning
Property prediction

ABSTRACT

Graph neural networks (GNNs) are emerging in chemical engineering for the end-to-end learning of physico-chemical properties based on molecular graphs. A key element of GNNs is the pooling function which combines atom feature vectors into molecular fingerprints. Most previous works use a standard pooling function to predict a variety of properties. However, unsuitable pooling functions can lead to unphysical GNNs that poorly generalize. We compare and select meaningful GNN pooling methods based on physical knowledge about the learned properties. The impact of physical pooling functions is demonstrated with molecular properties calculated from quantum mechanical computations. We also compare our results to the recent set2set pooling approach. We recommend using sum pooling for the prediction of properties that depend on molecular size and compare pooling functions for properties that are molecular size-independent. Overall, we show that the use of physical pooling functions significantly enhances generalization.

1. Introduction

Graph neural networks (GNNs) are emerging for end-to-end learning of molecular properties (Kearnes et al., 2016; Niepert et al., 2016; Hamilton et al., 2017a; Duvenaud et al., 2015; Rittig et al., 2022a; Reiser et al., 2022) in a broad variety of applications including chemical engineering (Schweidtmann et al., 2020; Li et al., 2021; Rittig et al., 2022b, 2023; Felton et al., 2022; Sanchez Medina et al., 2022a,b; Qin et al., 2022), (quantum) chemistry (Gilmer et al., 2017; Yang et al., 2019; Schütt et al., 2017; Wu et al., 2018; Back et al., 2019) and the prediction of physical (Coley et al., 2017) and crystal properties (Chen et al., 2019; Xie and Grossman, 2018). Although GNNs are flexible models for end-to-end learning, we show that their pooling function needs to be carefully selected because wrong decisions can lead to unphysical GNNs that are more prone to overfitting. GNNs take molecular graphs as inputs and represent atoms by nodes and bonds by edges. In addition, atoms and nodes are characterized by corresponding feature vectors. Most commonly, GNN architectures are based on message passing neural networks (MPNNs) (Gilmer et al., 2017). In MPNNs, the

node feature vectors are updated through a series of message passing procedures of neighboring nodes. Each of these sequential message passings corresponds to the graph convolutional layers of the GNN. Then, the resulting node feature vectors are combined into a molecular fingerprint vector through pooling. This molecular fingerprint is finally mapped to molecular properties of interest by feedforward artificial neural networks (ANNs). We show that using physical knowledge for the selection of the pooling function, which combines the feature vectors of all atoms into the molecular fingerprint, is critical for the GNN's performance.

Recent studies emphasize the importance of the choice of pooling functions in GNNs beyond molecular property prediction. These studies focus on the aggregation function of the neighborhood aggregation within the message-passing phase in GNNs, where information from neighboring nodes is aggregated and passed to a specific node. Xu et al. (2018a) examine sum, mean, and max pooling and conclude that sum pooling is more powerful than mean and max pooling since it can better distinguish different graph structures. Corso et al. (2020) generalize the

* Corresponding author at: RWTH Aachen University, Process Systems Engineering (AVT.SVT), Forckenbeckstr. 51, 52074 Aachen, Germany.
E-mail address: amitsos@alum.mit.edu (A. Mitsos).

sum aggregation and further propose to combine multiple aggregators, thereby outperforming GNN architectures with a single aggregator on various data sets. Even though these works focus on neighborhood aggregation, the theoretical implications are transferable to the pooling operation for obtaining a graph representation/molecular fingerprint vector since it is an aggregation of all node vectors within a graph. Hence, the selection of the pooling operators is of great importance for GNNs, with the sum pooling operator being favorable in most cases.

In molecular property prediction, the vast majority of previous works do not select pooling functions based on physical understanding. In the previous literature, common pooling functions are mean, sum, and max pooling (Wu et al., 2021). Most previous works use sum pooling (Xu et al., 2018a; Coley et al., 2017; Yang et al., 2019; Lu et al., 2019), while a few use mean pooling (Morris et al., 2019; Shindo and Matsumoto, 2019). Moreover, typically the same pooling function is used for a range of different properties. We argue that this can lead to unphysical GNNs and result in unnecessary errors. An illustrative example is the molecular weight that is given by the sum of the atom weights. In this case, using mean pooling in a standard GNN would lead to an unphysical architecture that cannot learn the correct underlying physics because the molecular mass cannot be computed as an average of atom weights. In contrast, selecting the sum pooling function according to the underlying physics enables the GNN to learn a meaningful model.

Some researchers circumvent the issue of selecting pooling functions by introducing flexible models for pooling such as set2set (Vinyals et al., 2015), DiffPool (Ying et al., 2018), or SortPool (Zhang et al., 2018). For example, the set2set approach employs a long short-term memory (LSTM) architecture designed for unordered and size-variant input sets (Vinyals et al., 2015). The authors of DiffPool propose a hierarchical GNN structure that progressively coarsens the input graph in each layer by aggregating clusters of nodes until a single graph representation is obtained (Ying et al., 2018). SortPool arranges the learned node representation in a consistently ordered tensor which is then truncated or extended to a user-defined fixed size (Zhang et al., 2018). Similarly, GNNs with a large number of convolutional layers combine node information through convolutions and thus reduce the importance of pooling functions (Morris et al., 2019). Advanced pooling methods have also been applied in molecular property prediction. For instance, Gilmer et al. (2017) applied the set2set method for learning various molecular properties from the QM9 data set (Ruddigkeit et al., 2012; Ramakrishnan et al., 2014) and achieved state-of-the-art accuracies on all target properties compared to other GNN models at the time of publication. However, the additional flexibility typically results in larger data requirements, higher model variance, and the risk of overfitting. In other words, the selection of physical pooling functions over flexible model architectures for pooling can be understood as enforcing a hybrid model structure, which is known to reduce the data demand (Psychogios and Ungar, 1992; Fiedler and Schuppert, 2008; Schweidtmann et al., 2021).

Only a few studies consider the physical nature of a property in the selection of the pooling function. Pronobis et al. (2018) state that decomposition of molecules into atom-wise contributions combined with a property-suitable pooling function works better for “extensive properties”. Other works use property-specific pooling functions, where mean/set2set or sum pooling is applied to “intensive” or “extensive” properties, respectively (Schütt et al., 2020, 2018; Gubaev et al., 2018; Ye et al., 2020; Liu et al., 2021). For example, Schütt et al. (2018) train a property prediction model on the QM9 data set and apply sum pooling for “extensive” properties and mean pooling for “intensive” properties (Schütt et al., 2018). Overall, however, the physical selection of pooling functions is somewhat contradictory in the previous literature and the terms “intensive” or “extensive” have been used colloquially and not in their thermodynamic sense. Also, a comparison of different pooling functions on the prediction and generalization capabilities of GNNs against a physical background has not been conducted yet, hence

there is no guide for selecting suitable pooling functions based on physical knowledge in the literature.

We evaluate GNN pooling functions against the underlying physical nature of the learned properties. Specifically, we analyze the impact of physical pooling functions on molecular properties learned from the common QM9 data set (Ruddigkeit et al., 2012; Ramakrishnan et al., 2014) and demonstrate their superior performance.

2. Materials and methods

We first describe the graph representation of molecules and then briefly introduce the general GNN architecture used for property prediction. Finally, we provide physical insight into the learned properties and use the insight to design physical GNN architectures.

2.1. Molecular graph

Molecules can be described as molecular graphs with nodes $v, w \in V$ representing atoms and edges $e_{vw} \in E$ representing bonds. Each atom is described by a feature vector $\mathbf{f}^V(v)$, containing atom information, e.g., atom mass or orbital hybridization. Similarly, each bond is described by a bond feature vector $\mathbf{f}^E(e_{vw})$ that contains information on the bond type, e.g., single or double bond. Commonly, for organic molecules the hydrogen (H) atoms are omitted and replaced by the hydrogen count as a node feature (Todeschini and Consonni, 2000); this results in reducing the complexity of molecular graphs and therefore reducing data demand.

2.2. Graph neural network

GNNs exhibit two phases (Gilmer et al., 2017) as shown in Fig. 1: (i) message passing phase and (ii) readout phase.

To initialize the message passing phase, each node $v \in V$ is assigned a state vector $\mathbf{h}_v^{l=0}$ initialized by the respective node feature vector (Gilmer et al., 2017). Then, the state vector of the nodes in layer l are updated with information from their neighboring nodes $w \in N(v)$ along edges e_{vw} :

$$\mathbf{h}_v^l = U_l \left(\mathbf{h}_v^{l-1}, \sum_{w \in N(v)} M_l(\mathbf{h}_v^{l-1}, \mathbf{h}_w^{l-1}, \mathbf{f}^E(e_{vw})) \right), \quad (1)$$

where $U_l(\cdot)$ and $M_l(\cdot)$ respectively denote the state update function and the message function in layer l . This message passing procedure is repeated L times until each node state vector \mathbf{h}_v^L includes information about its local environment. This iterative message passing corresponds to the stacking of L graph convolutional layers.

We consider a standard GNN including edge features in the message passing phase, also known as 1-GNN (Morris et al., 2019; Hamilton et al., 2017b). The 1-GNN uses the following message passing function:

$$\mathbf{h}_v^l = \sigma \left(\theta_v^l \cdot \mathbf{h}_v^{l-1} + \sum_{w \in N(v)} \text{ANN}_{\theta_e^l}(\mathbf{f}^E(e_{vw})) \cdot \mathbf{h}_w^{l-1} \right), \quad (2)$$

where σ indicates an activation function, θ_v^l denotes a parameter matrix, and $\text{ANN}_{\theta_e^l}$ denotes a feedforward ANN mapping the respective feature vectors $\mathbf{f}^E(e_{vw})$ of the edges e_{vw} connecting node v with its neighbors to a parameter matrix θ_e^l , referred to as edge feature network.

In the readout phase, the final state vectors of the nodes \mathbf{h}_v^L are combined into a graph state vector \mathbf{h}_G by a pooling function. The pooling function is necessary for molecular property prediction because the number of atoms usually differs between different molecules. This leads to a varying number of atom feature vectors that need to be combined to the molecular fingerprints. Thus, the pooling function combines a varying number of final state vectors for the nodes into a single graph state vector. In the context of molecular property prediction, the literature commonly refers to \mathbf{h}_G as the molecular fingerprint.

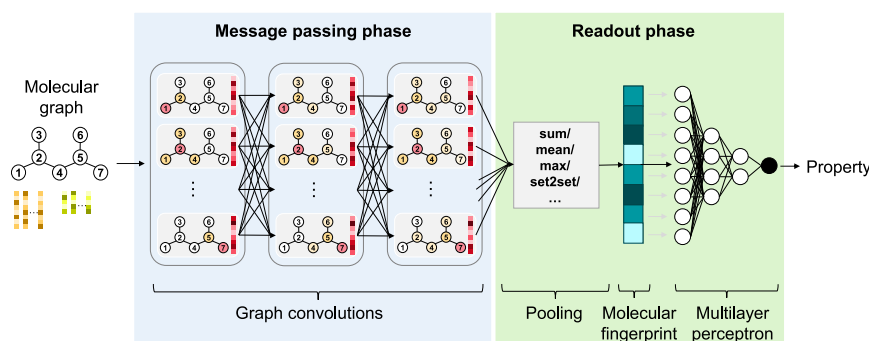


Fig. 1. Illustration of the GNN structure highlighting the message passing and readout phases.

This molecular fingerprint \mathbf{h}_G is finally fed into a feed-forward ANN for the prediction of molecular properties, $\hat{\mathbf{p}} = \text{MLP}(\mathbf{h}_G)$.

The molecular fingerprint is given by the pooling function $f_p(\cdot)$ that depends on the final state vectors of the nodes \mathbf{h}_v^L with $v \in V$:

$$\mathbf{h}_G = f_p(\{\mathbf{h}_v^L \mid v \in V\}) \quad (3)$$

Common choices for f_p are the sum, mean, and max functions. It is also possible to consider intermediate node state vectors instead of the final state vectors only in the pooling function, referred to as jumping knowledge (Xu et al., 2018b), which can lead to a beneficial structure-aware fingerprint vector. Most works, however, focus on pooling of the final state vectors. An alternative pooling function is the set2set method (Vinyals et al., 2015) which can capture more complex relationships between different atomic contributions (Gilmer et al., 2017; Schütt et al., 2018) by employing a long short-term memory (LSTM) model (Vinyals et al., 2015). After T steps of the following iterative computation, the molecular fingerprint is obtained by $\mathbf{h}_G = \mathbf{q}_T^*$ with

$$\begin{aligned} \mathbf{q}_t^* &= \mathbf{q}_t \parallel \mathbf{r}_t \\ \mathbf{q}_t &= \text{LSTM}(\mathbf{q}_{t-1}^*) \\ \mathbf{r}_t &= \sum_v \mathbf{a}_{v,t} \cdot \mathbf{h}_v^L \end{aligned} \quad (4)$$

$$\mathbf{a}_{v,t} = \text{softmax}(\mathbf{h}_v^L \cdot \mathbf{q}_t)$$

where \mathbf{q}_t is a query vector for iteration t providing information about the previous attention readout vector \mathbf{r}_t from the memories, $\mathbf{a}_{v,t}$ is an attention vector resulting from averaging the attention of a node v by applying the softmax function, \mathbf{r}_t is the attention readout, similar to the simple pooling method with sum, and \mathbf{q}_t is a concatenation (\parallel) of the current query vector and the attention readout. The vector \mathbf{q}_{t-1}^* is initialized at $t = 0$ by $\mathbf{q}_{-1}^* = \mathbf{0}$.

Recent GNNs incorporate physical knowledge into message passing. Over the last years, multiple MPNN architectures have been proposed that integrate physical knowledge to the message passing scheme, e.g., SchNet (Schütt et al., 2018), PhysNet (Unke and Muwly, 2019), DimeNet (Klicpera et al., 2020), MXMNet (Zhang et al., 2020a). This includes the incorporation of directional information, such as interatomic distances and angles between atom pairs, into the message function $M_l(\cdot)$ (cf. Eq. (1)) modeling the interactions of atoms. We consider MXMNet that utilizes physical-driven message passing while preserving computational efficiency (Zhang et al., 2020a). Within MXMNet, two message passing schemes are applied. In a global message passing scheme, information between atoms with a global cutoff distance d_g is exchanged. Further, a local message passing is applied to exchange information between atoms with a local cutoff distance d_l with $d_g > d_l$. This local cutoff distance represents the connectivity of atoms that are connected by chemical bonds. The architecture further enables to transfer of information between atom representations in the global and local message passing by including a cross layer mapping. In the readout step, the learned atom-wise representations are subsequently pooled by the sum operator for molecular property prediction.

2.3. Physical insight

The prediction of molecular properties by decomposing molecules into atomic contributions has a long history in chemical research. According to Bonchev and Rouvray (1991), the first investigations into properties of molecules with additive characteristics in terms of atomic contributions were carried out in the 1850s. Later, quantitative structure–property relationship (QSPR) and group additivity methods were developed based on the additive character of atoms or functional groups within a molecule (Katritzky et al., 1995; Benson et al., 1969; Gani et al., 1991; Joback and Reid, 1987). Yet, not every molecular property exhibits purely additive effects.

In thermodynamics, macroscopic properties are categorized as intensive or extensive (Cohen et al., 2007). A system property is extensive if it scales linearly with the mass (and as such “extent”) of the system; examples are the mass or volume. In contrast, intensive properties do not change with the system mass. Note that sometimes thermodynamicists also distinguish between intensive (e.g., temperature and pressure) and specific (extensive quantity divided by volume, e.g., density) properties (Stephan and Mayinger, 2013), but we will not. We transfer these concepts to molecules and distinguish between *molecular size-independent* and *molecular size-dependent* properties. Molecular size-dependent properties scale with the number of atoms in a molecule. For example, the molecular weight is determined by how many atoms of which type are present in a molecule. In contrast, there exist molecular size-independent properties that do not scale with the number of atoms in a molecule, e.g., the highest occupied molecular energy level (HOMO) (Schütt et al., 2018). Moreover, some properties of a substance, e.g., activity or toxicity, are mostly influenced by certain functional groups or structural fragments. Note that this dependency on molecular size does not necessarily correspond to the formal definition of intensive or extensive properties in a thermodynamic sense because the former is considered at a microscopic atom-based molecular level, not at a macroscopic mass-based system level. For example, the molar enthalpy with the unit J/mol is an intensive property. On a molecular level, the enthalpy of atomization $H_{298,atom}$ with the unit J/mol, is a molecular size-dependent property describing the amount of energy needed to break up a molecule into all of its single atoms at room temperature and fixed pressure (Gilmer et al., 2017).

Schütt et al. (2018) consider the QM9 properties dipole moment (μ), isotropic polarizability (α), electronic spatial extent (R^2), zero point vibrational energy (ZPVE), heat capacity at 298.15 K ($C_{v,298}$), atomization energy at 0 K ($U_{0,atom}$), atomization energy at 298.15 K ($U_{298,atom}$), enthalpy of atomization at 298.15 K ($H_{298,atom}$), free energy of atomization at 298.15 K ($G_{298,atom}$) as “extensive”. Other properties are the highest occupied molecular orbital (ϵ_{HOMO}), lowest unoccupied molecular orbital energy level (ϵ_{LUMO}), and HOMO-LUMO gap ($\Delta\epsilon$). For some of these properties, physical dependencies on molecular size are known. For example, Miller and Savchik (1979) developed a semi-empirical approach for the prediction of isotropic polarizabilities as a sum of atomistic contributions that depend on their hybridization

states based on theoretical calculations already in 1979. Each nonlinear molecule has $3N - 6$ vibrational degrees of freedom ($3N - 5$ for linear ones), N being its number of atoms. Each degree of freedom has a ZPVE proportional to its frequency $\nu \propto \sqrt{f/m}$. Here, m is the reduced mass of the parts of a molecule that vibrate with respect to each other and f is the force constant of this vibration. Hence, ZPVE is a molecular size-dependent property to first order but the frequencies of vibrations that include large fractions of a molecule decrease with increasing molecular size. This effect should be learned by the GNN. Similar relations apply to the heat capacity, enthalpy, and entropy contributions with the minor complication that terms dependent on the molecular mass (translation) and the moment of inertia (rotation) arise for some of the contributions (Atkins and Friedman, 2011). Atomization energies and enthalpies include essentially sums of contributions of all bonds, which may be non-local in the case of conjugated bonds, and are hence molecular size-dependent properties as well. The electronic spatial extent is determined mainly by the shapes of the orbitals in very small molecules and is closely related to the radius of gyration for large molecules, which may even depend on the solvent for large molecules such as polymers and may scale with a fractal exponent in this case.

The dipole moment (μ) is a particularly interesting property. Even though formally the molecule size enters the dipole moment definition equation, $\mu = \sum r_i q_i$, with the positions r_i of the charges q_i , in most molecules local functional groups determine μ , and depending on orientation they can even weaken each other. In particular, one or a few strongly polar groups dominate the dipole moment that can then be written as the sum of the individual dipole moment vectors. Thus, we classify the dipole moment as molecular size independent. The prediction of dipole moment is expected to be challenging for conventional GNNs because long-range orientational relations between the polar groups may need to be learned by a model, e. g., for describing the difference between the polar ortho- and the unpolar para-benzoquinone.

Similarly, the relations are complex for energies of the HOMO, the LUMO, and their difference (i.e., the HOMO-LUMO-gap). Depending on the type of molecule, these orbitals may be quite localized to a certain functional group and thus independent of molecular size in some molecules. However, they may also be delocalized in other molecules and thus dependent on molecular size for small and medium-sized molecules but converge to a limit for large molecules as can be seen, e.g., from the Hückel model for conjugated double bonds (Atkins and Friedman, 2011). Hence this property may be particularly challenging for a GNN model.

3. Results and discussion

In order to demonstrate the relevance of the pooling function in the readout phase, two case studies are conducted and discussed below. First, we consider the illustrative prediction of the molecular weight. Then, we consider the prediction of twelve quantum mechanical properties collected in the QM9 data set.

3.1. Hyperparameters and implementation

Our implementations are based on the models in PyTorch Geometric developed by Fey and Lenssen (2019). For our case study, we combine each mean, sum, and max as well as set2set pooling with the 1-GNN. The hyperparameters of the models are selected based on our experience from our previous work on predicting fuel properties (Schweidtmann et al., 2020). The molecular graphs have the following features encoded as a one-hot vector: (node) atom type, is aromatic, is in ring, hybridization (e.g., sp , sp^2 , sp^3) hydrogen count, (edge) bond type, conjugated, and stereo. The 1-GNN comprises three graph convolutional layers with hidden dimension size of 64 and ELU activation function. The graph convolutions further have an edge feature network (cf. Eq. (2)) with three layers with #1: 10, #2: 128,

#3: 4096 (64^2) layers and a ReLU activation function. To map the molecular fingerprint (\mathbf{h}_G) to the property (\hat{p}) of interest, we use multilayer perceptrons (MLPs), $\hat{p} = \text{MLP}(\mathbf{h}_G)$, with an ELU activation function. The MLPs constitute four layers with #1: 64, #2: 32 #3: 16, #4: 1 neurons when mean, sum, or max pooling is used. When the set2set pooling is used, the MLP layers have #1: 128, #2: 64, #3: 32, #4: 1 neurons, because the output vector of the set2set method is twice its input size. For the set2set method, we set the number of processing steps T to 3. For training, we use the ADAM optimizer, a mean squared error loss function, and a batch size of 128. The initial learning rate is set to 0.001 and is reduced by a factor of 0.7 during training if the validation loss has not decreased for 5 consecutive epochs.

We additionally test mean, sum, max pooling with MXMNet. We use the implementation and default hyperparameters as it was provided by the authors of MXMNet (Zhang et al., 2020a), cf. Zhang et al. (2020b): The network has six MXMNet module layers using a SiLU activation function and a global cutoff distance of 5. It is trained with the ADAM optimizer, a mean absolute error loss function, and a batch size of 128. The initial learning rate 0.0001 is decreased every epoch by a factor of 0.9961697.

3.2. Illustrative case study: Molecular weight

To illustrate the importance of physical pooling functions on a simple example, we learn the molecular weight of alkanes. To compose the data set, we obtain about 2300 alkanes from C_1H_4 to $C_{60}H_{122}$ from the PubChem database (Kim et al., 2016) and compute their molecular weight using RDKit (Landrum, 2020).

For illustration, we split this case study into two steps. Firstly, 1-GNNs with sum, mean, and max pooling functions are trained, validated, and tested on corresponding data sets with alkanes with up to 30 C-atoms. Secondly, the trained GNNs are tested against an external data set containing alkanes with more than 30 C-atoms. Thus, the generalization capabilities of the 1-GNNs are tested against extrapolated data.

The training of the 1-GNNs is repeated ten times for each pooling function with a maximum number of 500 epochs. The initial data set with up to 30 C-atoms is randomly split into 80% training, 10% validation, and 10% test sets for each run. Since we consider alkanes, we choose the attributes of the nodes in the molecular graph to include the hydrogen count only and do not use edge attributes, hence we replace the edge feature network in the message passing of the 1-GNN (cf. Eq. (2)) by a (learnable) parameter matrix (θ_w^l) that does not depend on the edge type.

Fig. 2 shows the test set performance of the 1-GNN with different pooling functions. Fig. 2(a) illustrates the test results for the data set of alkanes with up to 30 C-atoms. As expected, the sum pooling leads to the best performance on the test data set with an average mean absolute error of 0.06 g/mol as it captures the molecular size-dependent character of the molecular weights. In contrast, mean pooling leads to an average mean absolute error 0.7 g/mol. Max pooling even leads to an average absolute error in the order of 31 g/mol.

Fig. 2(b) shows the mean absolute error on the test data sets of alkanes with more than 30 C-atoms. The 1-GNN with sum pooling leads to an average absolute error of 6.5 g/mol. In contrast, the 1-GNN with mean pooling leads to an average error of 63 g/mol and the 1-GNN with max pooling leads to an average error of 213 g/mol.

The results clearly show that the sum pooling function, which was selected based on our physical insight, performs better than the unphysical mean and max pooling functions for the prediction of the molecular weight. In particular, the sum pooling outperforms the unphysical pooling functions significantly when extrapolating the model. These results support our theoretical expectations that GNNs with physical pooling functions lead to better predictive quality for both interpolation and extrapolation.




Fig. 2. Mean absolute error in g/mol for test of the 1-GNN with different pooling functions, namely: sum, mean, max. (a) test data set of alkanes with up to 30 C-atoms, (b) external data set of alkanes with 35 up to 60 C-atoms. Results are for ten independent training runs, each with 500 periods.

3.3. Physicochemical properties

We analyze the importance of physical pooling functions for a variety of relevant properties. In addition, we compare our results to a more complex set2set readout function and explore the MXMNet (Zhang et al., 2020a) architecture.

We use the QM9 data set to train our models (Ruddigkeit et al., 2012; Ramakrishnan et al., 2014). The experimental setup is twofold. First, the 1-GNNs with sum, mean, max, and set2set pooling functions are trained, validated, and tested on randomly selected subsets of the whole QM9 data set. This approach assesses the interpolation capabilities of the respective pooling functions. Second, we train, validate, and test the 1-GNN and MXMNet on the QM9 data excluding molecules with exactly 9 heavy atoms. These models are then tested against molecules with 9 heavy atoms from the QM9 data set. This approach is chosen to assess the generalization capabilities of the models in terms of extrapolation ability. For each training run, the data set is randomly split into 80% training, 10% validation, and 10% test sets. The training is stopped after 300 and 900 periods for the 1-GNN and MXMNet, respectively. The mean absolute error for the test set is reported based on the period with the lowest validation error.

3.3.1. Interpolation

The results for testing sum, mean, max, and set2set pooling function on the whole QM9 data set are summarized in Table 1. Overall, the results indicate that physically meaningful pooling functions lead to favorable performances on the QM9 data set for interpolation. It can be observed that for all molecular size-dependent properties, the 1-GNN with sum pooling significantly outperforms the 1-GNNs with mean and max pooling.

For the molecular size-independent properties, we do not observe a superior performance of one pooling function; all pooling functions result in similar accuracies, in most cases within the standard deviation of each other.

3.3.2. Generalization with 1-GNN architecture

In order to analyze the generalization capability of the pooling functions on the QM9 data set, we train the 1-GNNs only on molecules with up to 8 heavy atoms, i.e., a maximum number of 8 C, N, O, F atoms. Then, we test the prediction accuracy on an internal test set, i.e., containing molecules with up to 8 heavy atoms, and also on a data set with the remaining QM9 molecules that have exactly 9 heavy atoms. The latter test set therefore tests the extrapolation capability of the derived GNNs. The results of the extrapolation are summarized in Table 2.

The interpolation performance of the models (indicated in black in Table 2) is similar to that of the previous models trained on the whole QM9 data set (cf. Table 1). Notably, the absolute errors increase for

some properties, e.g., atomization energy, as the training set is much smaller. The QM9 data set contains about 108,000 molecules with 9 atoms and about 22,000 molecules with 1 to 8 heavy atoms.

For extrapolation (indicated in blue in Table 2), we find that sum pooling performs much better compared to mean, max, and set2set pooling on all tested molecular size-dependent properties. For the molecular size-independent properties we observe that the sum pooling does not outperform the other pooling functions anymore. Rather, mean and max pooling perform on a similar level or slightly better compared to sum pooling; mostly averaged accuracies within the standard deviation of the different pooling function are observed. Notably, the set2set method does not improve the accuracy compared to sum, mean, and max pooling for any property.

3.3.3. Generalization with MXMNet architecture

We also analyze the influence of the pooling function on the MXMNet GNN model (Zhang et al., 2020a). The MXMNet model reached state-of-the-art performance on several prediction tasks in QM9 (Zhang et al., 2020a). MXMNet includes directional information in its message passing process. For the readout step, sum pooling is applied in the original MXMNet model. We compare MXMNet performance with three different pooling functions: sum, mean, and max. Similar to the 1-GNN, we test MXMNet trained on molecules of QM9 with up to 8 heavy atoms against an internal test set and an external test set, i.e., extrapolating to molecules with 9 heavy atoms. The results are summarized in Table 3.

The interpolation performance of the MXMNet (indicated in black in Table 3) is highly favorably compared to the 1-GNN architecture, as expected. The MXMNet with sum pooling outperforms the other pooling approaches for all molecular size-dependent properties. This is in agreement with our expectations and previous observations on the 1-GNN architecture. For the molecular size-independent properties, we observe that sum, mean, and max pooling perform very similarly. Notably, the extrapolation performance of the MXMNet architecture (indicated in blue in Table 3) also significantly outperforms the 1-GNN architecture on all properties but R^2 .

The generalization results follow the same pattern as our previous observations. Again, we observe a significant advantage of sum pooling for all molecular size-dependent properties. For the molecular size-independent properties, we obtain similar performance of the pooling functions for LUMO. For μ , sum pooling performs only slightly better than mean and max pooling. For the HOMO and the HOMO-LUMO gap, however, mean and max pooling outperform sum pooling by a factor of more than 3. Here, sum pooling also has a considerably higher standard deviation. This demonstrates that also sum pooling can promote overfitting and instable training, thus preventing generalization in case of size-independent properties. Notably, the extrapolation error of the MXMNet is much larger for the unphysical pooling function compared to the extrapolation error of the simpler 1-GNN with the same pooling function. This indicates that the selection of physical pooling functions could be more important for more complex models.

Table 1

Mean absolute errors averaged over three training runs with the corresponding standard deviation for testing 1-GNN with different pooling functions, sum, mean, max, set2set against QM9 target properties for the interpolation case. Properties are categorized into molecular size-independent (“m. size-ind.”) and molecular size-dependent (“m. size-dep.”) according to Schütt et al. (2018). Errors of the best pooling function are bold type.

| Target | | | Pooling | | | |
|--------------|-------------------|-----------|----------------------------|---------------------|----------------------------|---------------------|
| | | | Sum | Mean | Max | set2set |
| m. size-dep. | α | a_0^3 | 0.301 \pm 0.008 | 0.469 \pm 0.013 | 0.482 \pm 0.011 | 0.583 \pm 0.238 |
| | R^2 | a_0^2 | 21.5 \pm 0.5 | 25.0 \pm 0.6 | 24.1 \pm 0.4 | 24.1 \pm 1.9 |
| | ZPVE | meV | 9.39 \pm 1.06 | 24.16 \pm 3.99 | 26.99 \pm 1.10 | 21.96 \pm 6.34 |
| | $C_{v,298}$ | cal/mol K | 0.149 \pm 0.032 | 0.203 \pm 0.011 | 0.197 \pm 0.007 | 0.197 \pm 0.050 |
| | $U_{0,atom}$ | eV | 0.117 \pm 0.005 | 0.345 \pm 0.018 | 0.357 \pm 0.010 | 0.437 \pm 0.181 |
| | $U_{298,atom}$ | eV | 0.117 \pm 0.005 | 0.363 \pm 0.045 | 0.366 \pm 0.021 | 0.386 \pm 0.119 |
| | $H_{298,atom}$ | eV | 0.123 \pm 0.007 | 0.329 \pm 0.035 | 0.383 \pm 0.026 | 0.390 \pm 0.149 |
| | $G_{298,atom}$ | eV | 0.112 \pm 0.001 | 0.293 \pm 0.017 | 0.328 \pm 0.016 | 0.313 \pm 0.087 |
| m. size-ind. | μ | Debye | 0.452 \pm 0.003 | 0.456 \pm 0.004 | 0.449 \pm 0.003 | 0.472 \pm 0.016 |
| | ϵ_{HOMO} | eV | 0.0928 \pm 0.0009 | 0.0931 \pm 0.0016 | 0.0926 \pm 0.0011 | 0.1337 \pm 0.0060 |
| | ϵ_{LUMO} | eV | 0.0931 \pm 0.0028 | 0.0937 \pm 0.0011 | 0.0938 \pm 0.0021 | 0.0937 \pm 0.0007 |
| | $\Delta\epsilon$ | eV | 0.1322 \pm 0.0019 | 0.1336 \pm 0.0008 | 0.1283 \pm 0.0003 | 0.1323 \pm 0.0032 |

Table 2

Mean absolute errors averaged over three independent training runs with the corresponding standard deviation for testing 1-GNN with different pooling functions, sum, mean, max, and set2set, against: (black) QM9 data set excluding molecules with 9 heavy atoms, (blue) only molecules with 9 heavy atoms of the QM9 data set. The errors of the best pooling function are bold type. Units are equivalent to those in Table 1.

| Target | 1-GNN | | | |
|-------------------|--|--|--|---|
| | Sum | Mean | Max | set2set |
| α | 0.302 \pm 0.012 1.385 \pm 0.369 | 0.816 \pm 0.079 8.445 \pm 0.064 | 0.727 \pm 0.114 8.654 \pm 0.405 | 1.164 \pm 0.390 7.835 \pm 0.727 |
| R^2 | 17.9 \pm 0.3 83.6 \pm 25.7 | 24.9 \pm 0.6 236.0 \pm 9.1 | 23.3 \pm 0.5 229.5 \pm 6.7 | 29.4 \pm 12.1 227.0 \pm 11.7 |
| ZPVE | 12.26 \pm 1.42 28.81 \pm 4.76 | 56.14 \pm 5.12 366.09 \pm 14.65 | 54.39 \pm 3.61 312.73 \pm 45.65 | 54.68 \pm 8.97 450.78 \pm 108.45 |
| $C_{v,298}$ | 0.158 \pm 0.004 0.650 \pm 0.315 | 0.312 \pm 0.004 3.509 \pm 0.101 | 0.291 \pm 0.014 3.360 \pm 0.407 | 0.309 \pm 0.087 3.129 \pm 0.048 |
| $U_{0,atom}$ | 0.180 \pm 0.018 1.365 \pm 1.073 | 0.816 \pm 0.025 8.082 \pm 0.339 | 0.766 \pm 0.117 8.100 \pm 0.474 | 0.773 \pm 0.161 9.613 \pm 3.044 |
| $U_{298,atom}$ | 0.171 \pm 0.012 1.239 \pm 0.937 | 0.759 \pm 0.028 8.368 \pm 0.229 | 0.757 \pm 0.104 8.117 \pm 0.561 | 0.631 \pm 0.052 8.220 \pm 0.200 |
| $H_{298,atom}$ | 0.181 \pm 0.011 1.197 \pm 0.899 | 0.724 \pm 0.053 8.275 \pm 0.405 | 0.802 \pm 0.085 8.125 \pm 0.277 | 0.706 \pm 0.130 10.424 \pm 1.989 |
| $G_{298,atom}$ | 0.167 \pm 0.014 1.142 \pm 0.689 | 0.690 \pm 0.060 7.671 \pm 0.148 | 0.705 \pm 0.063 7.363 \pm 0.035 | 0.757 \pm 0.198 8.751 \pm 2.596 |
| μ | 0.469 \pm 0.006 0.588 \pm 0.002 | 0.465 \pm 0.002 0.572 \pm 0.005 | 0.454 \pm 0.002 0.569 \pm 0.004 | 0.501 \pm 0.051 0.611 \pm 0.038 |
| ϵ_{HOMO} | 0.110 \pm 0.001 0.142 \pm 0.003 | 0.109 \pm 0.002 0.140 \pm 0.001 | 0.112 \pm 0.003 0.142 \pm 0.002 | 0.127 \pm 0.010 0.158 \pm 0.012 |
| ϵ_{LUMO} | 0.116 \pm 0.002 0.179 \pm 0.002 | 0.114 \pm 0.003 0.171 \pm 0.001 | 0.110 \pm 0.004 0.167 \pm 0.005 | 0.118 \pm 0.006 0.173 \pm 0.007 |
| $\Delta\epsilon$ | 0.161 \pm 0.002 0.226 \pm 0.004 | 0.156 \pm 0.001 0.219 \pm 0.002 | 0.150 \pm 0.004 0.208 \pm 0.004 | 0.156 \pm 0.005 0.219 \pm 0.007 |

4. Conclusion

GNNs have emerged as a promising deep learning technique for end-to-end molecular property prediction in chemical engineering. The selection of pooling functions in GNNs for property prediction should be based on physical knowledge because incorrect pooling functions can promote overfitting and weaken generalization. We identify the dependency of the learned property on the molecular size as key property for the selection of the pooling function: When a property is molecular size-dependent, the sum pooling function should be used. When a property is molecular size-independent, sum pooling can lead to poor generalization. We recommend to compare sum, mean, and max pooling functions for size-independent properties.

Table 3

Mean absolute errors averaged over three independent training runs with the corresponding standard deviation for testing MXMNet with different pooling functions, sum, mean, max, against: (black) QM9 data set excluding molecules with 9 heavy atoms, (blue) only molecules with 9 heavy atoms of the QM9 data set. Error of best pooling function are bold type. Units are equivalent to those in Table 1.

| Target | MXMNet | | |
|-------------------|--|--|--|
| | Sum | Mean | Max |
| α | 0.0781 \pm 0.0059 0.1887 \pm 0.0099 | 0.1331 \pm 0.0049 1.2853 \pm 0.1534 | 0.1347 \pm 0.0040 0.4569 \pm 0.0244 |
| R^2 | 1.78 \pm 0.07 104.84 \pm 9.49 | 2.55 \pm 0.04 199.26 \pm 14.02 | 2.77 \pm 0.26 182.34 \pm 4.87 |
| ZPVE | 1.44 \pm 0.01 2.26 \pm 0.03 | 3.70 \pm 0.40 44.92 \pm 9.35 | 2.78 \pm 0.22 8.07 \pm 0.07 |
| $C_{v,298}$ | 0.0325 \pm 0.0010 0.0891 \pm 0.0058 | 0.0503 \pm 0.0031 2.2484 \pm 0.6533 | 0.0515 \pm 0.0010 0.6094 \pm 0.2107 |
| $U_{0,atom}$ | 0.0111 \pm 0.0004 0.0265 \pm 0.0011 | 0.0728 \pm 0.0012 1.4095 \pm 0.5951 | 0.0710 \pm 0.0108 0.5047 \pm 0.1451 |
| $U_{298,atom}$ | 0.0114 \pm 0.0005 0.0265 \pm 0.0011 | 0.0729 \pm 0.0035 1.2072 \pm 0.4537 | 0.0720 \pm 0.0098 0.5025 \pm 0.1428 |
| $H_{298,atom}$ | 0.0115 \pm 0.0006 0.0265 \pm 0.0009 | 0.0723 \pm 0.0016 1.5855 \pm 0.9168 | 0.0724 \pm 0.0081 0.5039 \pm 0.1233 |
| $G_{298,atom}$ | 0.0122 \pm 0.0003 0.0271 \pm 0.0010 | 0.0662 \pm 0.0018 2.1216 \pm 1.0470 | 0.0677 \pm 0.0088 0.5009 \pm 0.1058 |
| μ | 0.0892 \pm 0.0027 0.1551 \pm 0.0066 | 0.0981 \pm 0.0015 0.1708 \pm 0.0048 | 0.1203 \pm 0.0049 0.1997 \pm 0.0091 |
| ϵ_{HOMO} | 0.0516 \pm 0.0004 0.2273 \pm 0.1028 | 0.0460 \pm 0.0008 0.0687 \pm 0.0013 | 0.0550 \pm 0.0028 0.0743 \pm 0.0024 |
| ϵ_{LUMO} | 0.0366 \pm 0.0012 0.0700 \pm 0.0012 | 0.0384 \pm 0.0001 0.0707 \pm 0.0002 | 0.0449 \pm 0.0006 0.0816 \pm 0.0027 |
| $\Delta\epsilon$ | 0.0766 \pm 0.0006 0.3896 \pm 0.1788 | 0.0729 \pm 0.0022 0.1192 \pm 0.0024 | 0.0783 \pm 0.0030 0.1263 \pm 0.0052 |

Our computational results support this hypothesis showing that physical GNN architectures generalize better than unphysical architectures. In future research, the physical selection of pooling functions should always be considered when predicting molecular properties with GNNs.

CRediT authorship contribution statement

Artur M. Schweidtmann: Conceptualization, Methodology, Validation, Formal analysis, Investigation, Writing – original draft, Writing – review & editing, Visualization. **Jan G. Rittig:** Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization. **Jana M. Weber:** Methodology, Validation, Formal analysis, Writing – review & editing. **Martin Grohe:** Methodology, Validation, Formal analysis,

Writing – review & editing. **Manuel Dahmen:** Methodology, Validation, Formal analysis, Writing – review & editing, Supervision. **Kai Leonhard:** Methodology, Validation, Formal analysis, Writing – review & editing. **Alexander Mitsos:** Methodology, Validation, Formal analysis, Resources, Writing – review & editing, Supervision, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data and software availability

Our implementations are based on the models in PyTorch Geometric developed by Fey and Lenssen (2019). The MXMNet implementation is provided by the authors of MXMNet (Zhang et al., 2020a), cf. Zhang et al. (2020b). We use the QM9 data set to train our models (Rudigkeit et al., 2012; Ramakrishnan et al., 2014). We provide the model implementations for this study via Zenodo at Schweidtmann et al. (2023).

Acknowledgments

Supported by the German Research Foundation (DFG) within the framework of the Excellence Strategy of the Federal Government and the Länder - Cluster of Excellence 2186 “The Fuel Science Center” (ID390919832). Simulations were performed with computing resources granted by RWTH Aachen University, Germany under projects thes0682 and rwth0731. MD received funding from the Helmholtz Association of German Research Centers. This project was also funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 466417970 – within the Priority Programme “SPP 2331: Machine Learning in Chemical Engineering”.

References

- Atkins, P., Friedman, R., 2011. *Molecular Quantum Mechanics*, fifth ed. Oxford University Press, Oxford.
- Back, S., Yoon, J., Tian, N., Zhong, W., Tran, K., Ulissi, Z.W., 2019. Convolutional neural network of atomic surface structures to predict binding energies for high-throughput screening of catalysts. *J. Phys. Chem. Lett.* 10 (15), 4401–4408.
- Benson, S.W., Cruickshank, F.R., Golden, D.M., Haugen, G.R., O’Neal, H.E., Rodgers, A.S., Shaw, R., Walsh, R., 1969. Additivity rules for the estimation of thermochemical properties. *Chem. Rev.* 69 (3), 279–324.
- Bonchev, D., Rouvray, D., 1991. *Chemical Graph Theory: Introduction and Fundamentals*. Gordon and Breach Science Publishers, New York, United States.
- Chen, C., Ye, W., Zuo, Y., Zheng, C., Ong, S.P., 2019. Graph networks as a universal machine learning framework for molecules and crystals. *Chem. Mater.* 31 (9), 3564–3572.
- Cohen, E.R., Cvitas, T., Frey, J.G., Holström, B., Kuchitsu, K., Marquardt, R., Mills, I., Pavese, F., Quack, M., Stohner, J., Strauss, H.L., Takami, M., Thor, A.J. (Eds.), 2007. *Quantities, Units and Symbols in Physical Chemistry*. The Royal Society of Chemistry, pp. 001–232.
- Coley, C.W., Barzilay, R., Green, W.H., Jaakkola, T.S., Jensen, K.F., 2017. Convolutional embedding of attributed molecular graphs for physical property prediction. *J. Chem. Inf. Model.* 57 (8), 1757–1772.
- Corso, G., Cavalleri, L., Beaini, D., Liò, P., Veličković, P., 2020. Principal neighbourhood aggregation for graph nets. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (Eds.), *Advances in Neural Information Processing Systems*, Vol.33. Curran Associates, Inc., pp. 13260–13271.
- Duvenaud, D.K., Maclaurin, D., Iparraguirre, J., Bombarell, R., Hirzel, T., Aspuru-Guzik, A., Adams, R.P., 2015. Convolutional networks on graphs for learning molecular fingerprints. In: Cortes, C., Lawrence, N.D., Lee, D.D., Sugiyama, M., Garnett, R. (Eds.), *Advances in Neural Information Processing Systems 28* (NIPS 2015). Curran Associates, Inc, pp. 2224–2232.
- Felton, K.C., Ben-Safar, H., Alexei, A.A., 2022. DeepGamma: A deep learning model for activity coefficient prediction. In: 1st Annual AAAI Workshop on AI To Accelerate Science and Engineering (AI2ASE).
- Fey, M., Lenssen, J.E., 2019. Fast graph representation learning with PyTorch Geometric. arXiv preprint arXiv:1903.02428v3, arXiv.
- Fiedler, B., Schuppert, A., 2008. Local identification of scalar hybrid models with tree structure. *IMA J. Appl. Math.* 73 (3), 449–476.
- Gani, R., Nielsen, B., Fredenslund, A., 1991. A group contribution approach to computer-aided molecular design. *AIChE J.* 37 (9), 1318–1332.
- Gilmer, J., Schoenholz, S.S., Riley, P.F., Vinyals, O., Dahl, G.E., 2017. Neural message passing for quantum chemistry. arXiv preprint arXiv:1704.01212v2, arXiv.
- Gubaev, K., Podryabinkin, E.V., Shapeev, A.V., 2018. Machine learning of molecular properties: locality and active learning. *J. Chem. Phys.* 148 (24), 241727.
- Hamilton, W., Ying, Z., Leskovec, J., 2017a. Inductive representation learning on large graphs. In: *Advances in Neural Information Processing Systems 30* (NIPS 2017). pp. 1024–1034.
- Hamilton, W.L., Ying, R., Leskovec, J., 2017b. Representation learning on graphs: Methods and applications. *IEEE Database Eng. Bull.* 40 (3), 52–74.
- Joback, K.G., Reid, R.C., 1987. Estimation of pure-component properties from group-contributions. *Chem. Eng. Commun.* 57 (1–6), 233–243.
- Katritzky, A.R., Lobanov, V.S., Karelson, M., 1995. QSPR: the correlation and quantitative prediction of chemical and physical properties from structure. *Chem. Soc. Rev.* 24 (4), 279.
- Kearnes, S., McCloskey, K., Berndl, M., Pande, V., Riley, P., 2016. Molecular graph convolutions: Moving beyond fingerprints. *J. Comput. Aided Mol. Des.* 30 (8), 595–608.
- Kim, S., Thiessen, P.A., Bolton, E.E., Chen, J., Fu, G., Gindulyte, A., Han, L., He, J., He, S., Shoemaker, B.A., Wang, J., Yu, B., Zhang, J., Bryant, S.H., 2016. PubChem substance and compound databases. *Nucleic Acids Res.* 44 (D1), D1202–13.
- Klicpera, J., Groß, J., Günnemann, S., 2020. Directional message passing for molecular graphs. arXiv preprint arXiv:2003.03123.
- Landrum, G., 2020. Rdkit: Open-source cheminformatics. accessed on 05.04.2020.
- Li, Y., Li, P., Yang, X., Hsieh, C.-Y., Zhang, S., Wang, X., Lu, R., Liu, H., Yao, X., 2021. Introducing block design in graph neural networks for molecular properties prediction. *Chem. Eng. J.* 414, 128817.
- Liu, Z., Lin, L., Jia, Q., Cheng, Z., Jiang, Y., Guo, Y., Ma, J., 2021. Transferable multilevel attention neural network for accurate prediction of quantum chemistry properties via multitask learning. *J. Chem. Inf. Model.* 61 (3), 1066–1082.
- Lu, C., Liu, Q., Wang, C., Huang, Z., Lin, P., He, L., 2019. Molecular property prediction: A multilevel quantum interactions modeling perspective. *Proc. AAAI Conf. Artif. Intell.* 33, 1052–1060.
- Miller, K.J., Savchik, J.A., 1979. A new empirical method to calculate average molecular polarizabilities. *J. Am. Chem. Soc.* 101, 7206.
- Morris, C., Ritzert, M., Fey, M., Hamilton, W., Lenssen, J., Rattan, G., Grohe, M., 2019. Weisfeiler and leman go neural: Higher-order graph neural networks. In: *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*, 27.01.-01.02.2019, Vol. 4602-4609. AAAI Press, Honolulu, Hawaii, United States.
- Niepert, M., Ahmed, M., Kutzkov, K., 2016. Learning convolutional neural networks for graphs. In: Balcan, M.F., Weinberger, K.Q. (Eds.), *Proceedings of the 33rd International Conference on Machine Learning*. In: *Proceedings of Machine Learning Research*, 48, PMLR, New York, New York, USA, pp. 2014–2023.
- Pronobis, W., Schütt, K.T., Tkatchenko, A., Müller, K.-R., 2018. Capturing intensive and extensive DFT/TDDFT molecular properties with machine learning. *Eur. Phys. J. B* 91 (8), 338.
- Psichogios, D.C., Ungar, L.H., 1992. A hybrid neural network-first principles approach to process modeling. *AIChE J.* 38 (10), 1499–1511.
- Qin, S., Jiang, S., Li, J., Balaprakash, P., van Lehn, R., Zavala, V., 2022. Capturing molecular interactions in graph neural networks: A case study in multi-component phase equilibrium. <http://dx.doi.org/10.26434/chemrxiv-2022-3tq4c>, ChemRxiv preprint.
- Ramakrishnan, R., Dral, P.O., Rupp, M., von Lilienfeld, O.A., 2014. Quantum chemistry structures and properties of 134 kilo molecules. *Sci. Data* 1, 140022, EP –.
- Reiser, P., Neubert, M., Eberhard, A., Torresi, L., Zhou, C., Shao, C., Metni, H., van Hoesel, C., Schopmans, H., Sommer, T., Friederich, P., 2022. Graph neural networks for materials science and chemistry. *Commun. Mater.* 3 (1), 93.
- Rittig, J.G., Ben Hicham, K., Schweidtmann, A.M., Dahmen, M., Mitsos, A., 2023. Graph neural networks for temperature-dependent activity coefficient prediction of solutes in ionic liquids. *Comput. Chem. Eng.* 171, 108153.
- Rittig, J.G., Gao, Q., Dahmen, M., Mitsos, A., Schweidtmann, A.M., 2022a. Graph neural networks for the prediction of molecular structure-property relationships. arXiv preprint arXiv:2208.04852.
- Rittig, J.G., Ritzert, M., Schweidtmann, A.M., Winkler, S., Weber, J.M., Morsch, P., Heufer, K.A., Grohe, M., Mitsos, A., Dahmen, M., 2022b. Graph machine learning for design of high-octane fuels. arXiv preprint arXiv:2206.00619.
- Ruddigkeit, L., van Deursen, R., Blum, L.C., Reymond, J.-L., 2012. Enumeration of 166 billion organic small molecules in the chemical universe database GDB-17. *J. Chem. Inf. Model.* 52 (11), 2864–2875.
- Sanchez Medina, E.I., Linke, S., Stoll, M., Sundmacher, K., 2022a. Gibbs-Helmholtz Graph Neural Network: capturing the temperature dependency of activity coefficients at infinite dilution. arXiv preprint arXiv:2212.01199.
- Sanchez Medina, E.I., Linke, S., Stoll, M., Sundmacher, K., 2022b. Graph Neural Networks for the prediction of infinite dilution activity coefficients. *Digit. Discov.* 1, 216–225.
- Schütt, K.T., Arbabzadah, F., Chmiela, S., Müller, K.R., Tkatchenko, A., 2017. Quantum-chemical insights from deep tensor neural networks. *Nature Commun.* 8 (1), 13890.

- Schütt, K.T., Saucedo, H.E., Kindermans, P.-J., Tkatchenko, A., Müller, K.-R., 2018. SchNet - a deep learning architecture for molecules and materials. *J. Chem. Phys.* 148 (24), 241722.
- Schütt, K.T., Tkatchenko, A., Müller, K.-R., 2020. Learning representations of molecules and materials with atomistic neural networks. In: Schütt, K.T., Chmiela, S., Liliensfeld, O.A.v., Tkatchenko, A., Tsuda, K., Müller, K.-R. (Eds.), *Machine Learning Meets Quantum Physics*. In: *Lecture Notes in Physics*, 968, Springer International Publishing, Cham, pp. 215–230.
- Schweidtmann, A.M., Esche, E., Fischer, A., Kloft, M., Repke, J.-U., Sager, S., Mitsos, A., 2021. Machine learning in chemical engineering: A perspective. *Chem. Ing. Tech.*
- Schweidtmann, A.M., Rittig, J.G., König, A., Grohe, M., Mitsos, A., Dahmen, M., 2020. Graph neural networks for prediction of fuel ignition quality. *Energy Fuels* 34 (9), 11395–11407.
- Schweidtmann, A., Rittig, J.G., Weber, J.M., Grohe, M., Dahmen, M., Leonhard, K., Mitsos, A., 2023. Software for “Physical Pooling Functions in Graph Neural Networks for Molecular Property Prediction”. Zenodo, <http://dx.doi.org/10.5281/zenodo.7555594>.
- Shindo, H., Matsumoto, Y., 2019. Gated graph recursive neural networks for molecular property prediction. *arXiv preprint arXiv:1909.00259*.
- Stephan, K., Mayinger, F., 2013. *Thermodynamik. Grundlagen Und Technische Anwendungen: Band 1: Einstoffsysteme, Vol. 1*. Springer-Verlag.
- Todeschini, R., Consonni, V., 2000. *Handbook of Molecular Descriptors*. Wiley-VCH Verlag GmbH, Weinheim, Germany.
- Unke, O.T., Meuwly, M., 2019. PhysNet: A neural network for predicting energies, forces, dipole moments, and partial charges. *J. Chem. Theory Comput.* 15 (6), 3678–3693.
- Vinyals, O., Bengio, S., Kudlur, M., 2015. Order Matters: Sequence to sequence for sets. *arXiv preprint arXiv:1511.06391*.
- Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., Yu, P.S., 2021. A comprehensive survey on graph neural networks. *IEEE Trans. Neural Netw. Learn. Syst.* 32 (1), 4–24.
- Wu, Z., Ramsundar, B., Feinberg, E.N., Gomes, J., Geniesse, C., Pappu, A.S., Leswing, K., Pande, V., 2018. MoleculeNet: A benchmark for molecular machine learning. *Chem. Sci.* 9 (2), 513–530.
- Xie, T., Grossman, J.C., 2018. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Phys. Rev. Lett.* 120 (14), 145301.
- Xu, K., Hu, W., Leskovec, J., Jegelka, S., 2018a. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*.
- Xu, K., Li, C., Tian, Y., Sonobe, T., Kawarabayashi, K.-i., Jegelka, S., 2018b. Representation learning on graphs with jumping knowledge networks. In: Dy, J., Krause, A. (Eds.), *Proceedings of the 35th International Conference on Machine Learning*. In: *Proceedings of Machine Learning Research*, 80, PMLR, pp. 5453–5462.
- Yang, K., Swanson, K., Jin, W., Coley, C., Eiden, P., Gao, H., Guzman-Perez, A., Hopper, T., Kelley, B., Mathea, M., Palmer, A., Settels, V., Jaakkola, T., Jensen, K., Barzilay, R., 2019. Analyzing learned molecular representations for property prediction. *J. Chem. Inf. Model.* 59 (8), 3370–3388.
- Ye, S., Liang, J., Liu, R., Zhu, X., 2020. Symmetrical graph neural network for quantum chemistry with dual real and momenta space. *J. Phys. Chem. A* 124 (34), 6945–6953.
- Ying, Z., You, J., Morris, C., Ren, X., Hamilton, W., Leskovec, J., 2018. Hierarchical graph representation learning with differentiable pooling. In: *Advances in Neural Information Processing Systems*. pp. 4800–4810.
- Zhang, M., Cui, Z., Neumann, M., Chen, Y., 2018. An end-to-end deep learning architecture for graph classification. *Proc. AAAI Conf. Artif. Intell.* 32 (1).
- Zhang, S., Liu, Y., Xie, L., 2020a. Molecular mechanics-driven graph neural network with multiplex graph for molecular structures. In: *NeurIPS-W*.
- Zhang, S., Liu, Y., Xie, L., 2020b. Molecular mechanics-driven graph neural network with multiplex graph for molecular structures. accessed on 10.06.2021 <https://github.com/zetayue/MXMNet>.