

Sparse hierarchical representation learning on molecular graphs

Matthias Bal*
GTN LTD
London, UK
matthias.bal@gtn.ai

Hagen Triendl*
GTN LTD
London, UK
hagen.triendl@gtn.ai

Mariana Assmann
GTN LTD
London, UK
mariana.assmann@gtn.ai

Michael Craig
GTN LTD
London, UK
michael.craig@gtn.ai

Lawrence Phillips
GTN LTD
London, UK
lawrence.phillips@gtn.ai

Jarvist Moore Frost
GTN LTD
London, UK
jarvist.frost@gtn.ai

Usman Bashir
GTN LTD
London, UK
usman.bashir@gtn.ai

Noor Shaker
GTN LTD
London, UK
noor@gtn.ai

Vid Stojevic
GTN LTD
London, UK
vstojevic@gtn.ai

ABSTRACT

Architectures for sparse hierarchical representation learning have recently been proposed for graph-structured data, but so far assume the absence of edge features in the graph. We close this gap and propose a method to pool graphs with edge features, inspired by the hierarchical nature of chemistry. In particular, we introduce two types of pooling layers compatible with an edge-feature graph-convolutional architecture and investigate their performance for molecules relevant to drug discovery on a set of two classification and two regression benchmark datasets of MoleculeNet. We find that our models significantly outperform previous benchmarks on three of the datasets and reach state-of-the-art results on the fourth benchmark, with pooling improving performance for three out of four tasks, keeping performance stable on the fourth task, and generally speeding up the training process.

CCS CONCEPTS

• **Neural Networks** → *Statistical relational learning*; • **Physical sciences and engineering** → *Chemistry*.

KEYWORDS

Graph Neural Networks, Pooling, Molecular Graph

ACM Reference Format:

Matthias Bal, Hagen Triendl, Mariana Assmann, Michael Craig, Lawrence Phillips, Jarvist Moore Frost, Usman Bashir, Noor Shaker, and Vid Stojevic. 2019. Sparse hierarchical representation learning on molecular graphs. In *DLG '19: Deep Learning on Graphs, August 04–08, 2019, Anchorage, AK*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

*Both authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

DLG '19, August 04–08, 2019, Anchorage, AK

© 2019 Association for Computing Machinery.
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM... \$15.00
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION AND RELATED WORK

Predicting chemical properties of molecules has become a prominent application of neural networks in recent years. A standard approach in chemistry is to conceptualize groups of individual atoms as functional groups with characteristic properties, and infer the properties of a molecule from a multi-level understanding of the interactions between functional groups. This approach reflects the hierarchical nature of the underlying physics and can be formally understood in terms of renormalization [23]. It thus seems natural to use machine learning models that learn graph representations of chemical space in a local and hierarchical manner. This can be realized by coarse-graining the molecular graph in a step-wise fashion, with nodes representing effective objects such as functional groups or rings, connected by effective interactions.

Much published work leverages node locality by using graph-convolutional networks with message passing to process local information, see Gilmer et al. [15] for an overview. In graph classification and regression tasks, usually, only a global pooling step is applied to aggregate node features into a feature vector for the entire graph [8, 10, 15, 22].¹

An alternative is to aggregate node representations into clusters, which are then represented by a coarser graph [6, 9, 12, 25–27, 30]. Early work uses predefined and fixed cluster assignments during training, obtained by a clustering algorithm applied to the input graph. More recently, dynamic cluster assignments are made on learned node features [7, 13, 14, 34]. A pioneering step in using learnable parameters to cluster and reduce the graph was the DIFFPOOL layer introduced by Ying et al. [34]. Unfortunately, DIFFPOOL is tied to a disadvantageous quadratic memory complexity that limits the size of graphs and cannot be used for large sparse graphs. A sparse, and therefore more efficient, technique has been proposed by Gao and Ji [14] and further tested and explored by Cangea et al. [7], Gao et al. [13].

¹In some publications [1, 19] the phrase “pooling layer” has been used to refer to a MAX aggregation step. We reserve the notion of pooling for an operation which creates a true hierarchy of graphs, in line with its usage for images in computer vision.

Sparse pooling layers have so far not been developed for networks on graphs with both node and edge features. This is particularly important for molecular datasets, where edge features may describe different bond types or distances between atoms. When coarsening the molecular graph, new, effective edges need to be created whose edge features represent the effective interactions between the effective nodes. In this paper we explore two types of sparse hierarchical representation learning methods for molecules that process edge features differently during pooling: a *simple pooling* layer simply aggregates the features of the involved edges, while a more physically-inspired *coarse-graining* pooling layer determines the effective edge features using neural networks.

We evaluate our approach on established molecular benchmark datasets [33], in particular on the regression datasets ESOL and lipophilicity and the classification datasets BBBP and HIV, on which various models have been benchmarked [2–5, 11, 16, 18, 20, 29, 31, 32, 35]. We obtain significantly better results on the datasets ESOL, lipophilicity, and BBBP, and obtain state-of-the-art results on HIV. Simple pooling layers improve results on BBBP and HIV, while coarse-grain pooling improves results on lipophilicity. In general pooling layers can keep performance at least stable while speeding up training.

2 APPROACH

2.1 Model architecture

We represent input graphs in a sparse representation using node (\mathbf{a}) and edge (\mathbf{e}) feature vectors

$$\mathbf{a}_i^{(0)} = \mathbf{a}_i, \quad i = 1, \dots, n_{\text{nodes}}, \quad (1)$$

$$\mathbf{e}_{ij}^{(0)} = \mathbf{e}_{ij}, \quad i, j = 1, \dots, n_{\text{nodes}} \text{ for } j \in \text{NN}(i), \quad (2)$$

where j belongs to the set of nearest-neighbours (NN) of i . For chemical graphs we encode the atom type as a one-hot vector and its node degree as an additional entry in \mathbf{a}_i , while the bond type is one-hot encoded in \mathbf{e}_{ij} . Framed in the message-passing framework [15], the graph-convolutional models we use consist of alternating message-passing steps to process information locally and pooling steps that reduce the graph to a simpler sub-graph. Finally, a read-out phase gathers the node features and computes a feature vector for the whole graph that is fed through a simple perceptron layer in the final prediction step.

Dual-message graph-convolutional layer Since edge features are an important part of molecular graphs, the model architecture is chosen to give more prominence to edge features. We design a *dual-message graph-convolutional layer* that supports both node and edge features and treats them similarly. First, we compute an aggregate message $\mathbf{m}_i^{(k+1)}$ to a target node from all neighbouring source nodes $j \in \text{NN}(i)$ using a fully-connected neural network $f_{\mathbf{W}_a}$ acting on the source node features $\mathbf{a}_i^{(k)}$ and the edge features $\mathbf{e}_{ij}^{(k)}$ of the connecting edge. A self-message $\mathbf{s}_i^{(k+1)} = \mathbf{W}_s^{(k)} \mathbf{a}_i^{(k)} + \mathbf{b}_s^{(k)}$ from the original node features is added to the aggregated result. New node features are computed by applying batch norm (BN) and

a ReLU non-linearity, i.e.

$$\mathbf{m}_i^{(k+1)} = \sum_{j \in \text{NN}(i)} f_{\mathbf{W}_a} \left(\mathbf{a}_j^{(k)}, \mathbf{e}_{ij}^{(k)} \right), \quad (3)$$

$$\tilde{\mathbf{a}}_i^{(k+1)} = \text{ReLU} \left(\text{BN} \left(\mathbf{m}_i^{(k+1)} + \mathbf{s}_i^{(k+1)} \right) \right). \quad (4)$$

In contrast to the pair-message graph-convolutional layer of Gilmer et al. [15], we also update the edge feature with the closest node feature vectors via

$$\mathbf{m}_{ij}^{(k+1)} = g_{\mathbf{W}_e} \left(\tilde{\mathbf{a}}_i^{(k+1)} + \tilde{\mathbf{a}}_j^{(k+1)}, \mathbf{e}_{ij}^{(k)} \right), \quad (5)$$

$$\tilde{\mathbf{e}}_{ij}^{(k+1)} = \text{ReLU} \left(\text{BN} \left(\mathbf{m}_{ij}^{(k+1)} + \mathbf{s}_{ij}^{(k)} \right) \right), \quad (6)$$

where g is a fully-connected neural network and $\mathbf{s}_{ij}^{(k)} = \mathbf{W}_e^{(k)} \mathbf{e}_{ij}^{(k)} + \mathbf{b}_e^{(k)}$ is the edge feature self-message.

Pooling layer Pooling layers, as introduced in Gao and Ji [14], reduce the number of nodes by a fraction

$$\rho = K/n_{\text{nodes}}^{(k)}, \quad (7)$$

specified as a hyperparameter, via scoring all nodes using a learnable projection vector $\mathbf{p}^{(k)}$, and then selecting the K nodes with the highest score $y_i^{(k)}$. In order to make the projection vector $\mathbf{p}^{(k)}$ trainable, and thus the node selection differentiable, $\mathbf{p}^{(k)}$ is also used to determine a gating for each feature vector via

$$y_i^{(k)} = \mathbf{p}^{(k)} \cdot \tilde{\mathbf{a}}_i^{(k)}, \quad \mathbf{a}_i^{(k)} = \tilde{\mathbf{a}}_i^{(k)} \tanh \left(y_i^{(k)} \right), \quad (8)$$

where we only keep the top- K nodes and their gated feature vectors $\mathbf{a}_i^{(k)}$.

Pooling nodes requires the creation of new, effective edges between kept nodes while keeping the graph sparse. We discuss in Section 2.2 how to solve this problem in the presence of edge features.

Gather layer After graph-convolutional and pooling layers, a graph gathering layer is required to map from node and edge features to a global feature vector. Assuming that the dual-message message-passing steps are powerful enough to distribute the information contained in the edge features to the adjacent node features, we gather over node features only by concatenating MAX and SUM, and acting with a tanh non-linearity on the result. All models have an additional linear layer that acts on each node individually before applying the gather layer and a final perceptron layer.

2.2 Pooling with edge features

An important step of the pooling process is to create new edges based on the connectivity of the nodes before pooling in order to keep the graph sufficiently connected. For graphs with edge features this process also has to create new edge features. In addition, the algorithm must be parallel for performance reasons.

We tackle these issues by specifying how to combine edge features into an effective edge feature between remaining (kept) nodes. If a single dropped node or a pair of dropped nodes connect two kept nodes, we construct a new edge and drop the the ones linked to the dropped nodes. (see Fig. 1).

We propose two layers to calculate the replacement effective edge feature from the dropped edge features. A *simple pooling* layer computes an effective edge-feature by summing all edge feature

Model	RMSE results on		ROC-AUC results on	
	ESOL	Lipophilicity	BBBP	HIV
RF	1.07 ± 0.19	0.876 ± 0.040	0.714 ± 0.000	—
Multitask	1.12 ± 0.15	0.859 ± 0.013	0.688 ± 0.005	0.698 ± 0.037
XGBoost	0.912 ± 0.000 ^a	0.799 ± 0.054	0.696 ± 0.000	0.756 ± 0.000
KRR	1.53 ± 0.06	0.899 ± 0.043	—	—
GC	0.648 ± 0.019 ^a	0.655 ± 0.036	0.690 ± 0.009	0.763 ± 0.016
DAG	0.82 ± 0.08	0.835 ± 0.039	—	—
Weave	0.553 ± 0.035 ^a	0.715 ± 0.035	0.671 ± 0.014	0.703 ± 0.039
MPNN	0.58 ± 0.03	0.719 ± 0.031	—	—
Logreg	—	—	0.699 ± 0.002	0.702 ± 0.018
KernelSVM	—	—	0.729 ± 0.000	0.792 ± 0.000
IRV	—	—	0.700 ± 0.000	0.737 ± 0.000
Bypass	—	—	0.702 ± 0.006	0.693 ± 0.026
Chemception [5, 16]	—	—	—	0.752
Smiles2vec [2]	0.63	—	—	0.8
ChemNet [4]	—	—	—	0.8
Dummy super node GC [20]	—	—	—	0.766
EAGCN [29]	—	0.61 ± 0.02	—	0.83 ± 0.01
Mol2vec [18]	0.79	—	—	—
Outer RNN [31]	0.62	0.64	—	—
PotentialNet [11]	0.490 ± 0.014	—	—	—
SA-BILSTM [35]	—	—	—	0.83 ± 0.02
RNN encoder [32]	0.58	0.62	0.74	—
NoPool	0.410 ± 0.023	0.551 ± 0.010	0.846 ± 0.011	0.825 ± 0.008
SimplePooling (0.9)	0.410 ± 0.018	0.536 ± 0.009	0.839 ± 0.022	0.824 ± 0.014
SimplePooling (0.8)	0.417 ± 0.027	0.542 ± 0.013	0.869 ± 0.010	0.816 ± 0.020
SimplePooling (0.7)	0.485 ± 0.020	0.563 ± 0.016	0.859 ± 0.009	0.825 ± 0.015
SimplePooling (0.6)	0.413 ± 0.021	0.622 ± 0.030	0.852 ± 0.006	0.840 ± 0.019
SimplePooling (0.5)	0.437 ± 0.016	0.637 ± 0.027	0.851 ± 0.012	0.822 ± 0.019
CoarseGrainPooling (0.9)	0.420 ± 0.015	0.517 ± 0.005	0.852 ± 0.010	0.834 ± 0.015
CoarseGrainPooling (0.8)	0.430 ± 0.019	0.529 ± 0.020	0.853 ± 0.009	0.833 ± 0.009
CoarseGrainPooling (0.7)	0.472 ± 0.013	0.530 ± 0.005	0.856 ± 0.012	0.830 ± 0.007
CoarseGrainPooling (0.6)	0.495 ± 0.053	0.536 ± 0.026	0.838 ± 0.020	0.824 ± 0.026
CoarseGrainPooling (0.5)	0.412 ± 0.031	0.535 ± 0.009	0.858 ± 0.023	0.826 ± 0.010

Table 1: (Top) Literature results for the MoleculeNet benchmarks comparing RMSE and ROC-AUC results, on a range of models. Benchmarks without reference come from Wu et al. [33], except those values decorated with ^a, which come from Feinberg et al. [11]. (Bottom) Our model with coarse-grain pooling, simple pooling, and without pooling. The number in brackets specifies the pooling keep ratio ρ of the pooling layer.

vectors along the paths connecting pairs of kept nodes. When multiple paths between a pair of nodes are simultaneously reduced, this method will generate overlapping effective edge features. We reduce these to a single vector of the sum of overlapping edge feature vectors.

We know however that in chemistry effective interactions are more complex functions of the involved component features. Using this as an inspiration, we propose a more expressive *coarse-graining pooling* layer, which is obtained by replacing the simple aggregation function with neural networks to compute effective edge features. In particular, we use two fully-connected neural networks. The first network maps the atom and adjoining edge feature vectors of dropped nodes to a single effective-edge feature. The second network calculates effective edge features for kept edges (between kept nodes) to account for an effective coarse-grained interaction compensating for deleted nodes.

We use pooling layers after every convolutional layer except for the final one. For N convolutional layers, the number of nodes thus gets reduced by a factor ρ^{N-1} . This compression not only gets rid of irrelevant information but also reduces memory requirements and makes training faster, as we show in the experiments in Sec. 3.

3 EXPERIMENTAL RESULTS ON MOLECULENET

Model parameters and implementation We use hyperparameter tuning with the hyperband algorithm [21] to decide on the number of stacks and channel dimensions of graph-convolutional and pooling layers while keeping the pooling keep ratio defined in equation 7 fixed. All our models were implemented in PyTorch and trained on a single Nvidia Tesla K80 GPU using the ADAM optimizer with a learning rate of 0.0001.




Figure 1: Schematic of a graph pooling step (yellow nodes are kept, blue nodes are dropped). Dangling nodes are removed, together with their connecting edges. Pairs of edges connecting a dropped node to two kept nodes are coarse-grained to a new edge (heavy lines). New edges can also be constructed between kept nodes connected by two dropped nodes (heaviest line).

Evaluation on MoleculeNet We evaluate our models with and without pooling layers on the MoleculeNet benchmark set [33]. We focus on four different datasets, comprised of the regression benchmarks ESOL (1128 molecules) and Lipophilicity (4200 molecules), where performance was evaluated by RMSE, and the classification benchmarks on the BBBP (2039 molecules) and HIV (41127 molecules) datasets, evaluated via ROC-AUC. Following Wu et al. [33], we used a scaffold split for the classification datasets as provided by the DeepChem package. Apart from the benchmarks generated in the original paper, various models have been evaluated on these datasets [2–5, 11, 16, 18, 20, 29, 31, 32, 35]. An overview of the results in the literature can be found in the top of Table 1. Our results are the mean and standard deviation of 5 runs over 5 random splits (ESOL, Lipophilicity) or 5 runs over the same scaffold split (BBBP, HIV). Datasets were split into training (80%), validation (10%), and held-out test sets (10%). The validation set was used to tune model hyperparameters. All reported metrics are results on the test set. The results of our models with and without pooling are displayed in the lower part of the table.

pooling keep ratio	0.9	0.8	0.7	0.6	0.5
Speed-up	16%	24%	47%	55%	70%

Table 2: Speed-up of pooling runs of the HIV data set using Simple Pooling. The speed-up is measured as increase in speed in terms of elapsed real time compared to the run without pooling layer.

For the regression tasks, we found that our models significantly outperformed previous models for both datasets, with pooling layers keeping performance stable for ESOL and the coarse-grain pooling layer significantly improving results for Lipophilicity (see Table 1). Regarding classification tasks, we found that our models significantly outperformed previous models on BBBP and also exceeded previous benchmarks for the HIV dataset. For both datasets simple pooling layers improved performance. Curiously, the extent to which pooling layers improve performance and which layer is better suited for a particular task strongly depends on the dataset. It seems that simple pooling performs much better for classification tasks while for regression tasks it depends on the dataset.

We also measure the speed-up given by pooling layers during the evaluation on the HIV dataset in terms of elapsed real-time, using the simple pooling layer. The results are displayed in Table 2. We see significant speed-ups for moderate values of the pooling ratio.

4 CONCLUSION

We introduce two graph-pooling layers for sparse graphs with node and edge features and evaluate their performance on molecular graphs. While our model without pooling significantly outperforms benchmarks on ESOL, lipophilicity and BBBP and reaches state-of-the-art results on HIV in the MoleculeNet dataset, we find that our pooling methods improve performance and provide a speedup of up to 70% in the training of graph-convolutional neural networks that utilize edge features, along with a reduction in memory requirements.

While all experiments have been performed on datasets comprised of small, druglike molecules, we expect even stronger performance for datasets comprised of larger graphs like protein structures, where pooling can create a large, sequential hierarchy of graphs. More generally, our work may result in more pertinent and information-effective latent space representations for graph-based machine learning models.

REFERENCES

- [1] Han Altae-Tran, Bharath Ramsundar, Aneesh S. Pappu, and Vijay S. Pande. 2017. Low Data Drug Discovery with One-shot Learning. *American Chemical Society* 3 (2017). <http://arxiv.org/abs/1611.03199>
- [2] Garrett B. Goh, Nathan O. Hodas, Charles Siegel, and Abhinav Vishnu. 2017. SMILES2Vec: An Interpretable General-Purpose Deep Neural Network for Predicting Chemical Properties. (12 2017). arXiv:<https://arxiv.org/abs/1712.02034>
- [3] Garrett B. Goh, Charles Siegel, Abhinav Vishnu, and Nathan Hodas. 2018. Using Rule-Based Labels for Weak Supervised Learning: A ChemNet for Transferable Chemical Property Prediction. 302–310. <https://doi.org/10.1145/3219819.3219838>
- [4] Garrett B. Goh, Charles Siegel, Abhinav Vishnu, and Nathan O. Hodas. 2017. ChemNet: A Transferable and Generalizable Deep Neural Network for Small-Molecule Property Prediction. (12 2017). arXiv:<https://arxiv.org/abs/1712.02734>
- [5] Garrett B. Goh, Charles Siegel, Abhinav Vishnu, Nathan O. Hodas, and Nathan Baker. 2017. Chemception: A Deep Neural Network with Minimal Chemistry Knowledge Matches the Performance of Expert-developed QSARQSPR Models. (06 2017). arXiv:<https://arxiv.org/abs/1706.06689> <https://arxiv.org/abs/1706.06689>
- [6] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. 2013. Spectral Networks and Locally Connected Networks on Graphs. *preprint* (2013). arXiv:1312.6203 <http://arxiv.org/abs/1312.6203>
- [7] CAĆtĂclina Cangea, Petar Velićković, Nikola Jovanović, Thomas Kipf, and Pietro Liš. 2018. Towards Sparse Hierarchical Graph Classifiers. *Workshop on Relational Representation Learning (R2L) at NIPS* (2018). arXiv:1811.01287 <http://arxiv.org/abs/1811.01287>
- [8] Hanjun Dai, Bo Dai, and Le Song. 2016. Discriminative Embeddings of Latent Variable Models for Structured Data. *Proceedings of the International Conference on Machine Learning (ICML)* 48 (2016). arXiv:1603.05629 <http://arxiv.org/abs/1603.05629>
- [9] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. 2016. Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering. *Advances in Neural Information Processing Systems (NIPS)* 29 (2016). arXiv:1606.09375 <http://arxiv.org/abs/1606.09375>
- [10] David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P Adams. 2015. Convolutional Networks on Graphs for Learning Molecular Fingerprints. In *Advances in Neural Information Processing Systems* 28, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (Eds.). Curran Associates, Inc., 2224–2232. <https://arxiv.org/abs/1509.09292>
- [11] Evan N. Feinberg, Debnil Sur, Zhenqin Wu, Brooke E. Husic, Huanghao Mai, Yang Li, Saisai Sun, Jianyi Yang, Bharath Ramsundar, and Vijay S. Pande. 2018. PotentialNet for Molecular Property Prediction. *ACS Central Science* 4, 11 (2018), 1520–1530. <https://doi.org/10.1021/acscentsci.8b00507>

- arXiv:https://doi.org/10.1021/acscentsci.8b00507
- [12] Matthias Fey, Jan Eric Lenssen, Frank Weichert, and Heinrich Müller. 2018. SplineCNN: Fast Geometric Deep Learning with Continuous B-Spline Kernels. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2018). arXiv:1711.08920 <http://arxiv.org/abs/1711.08920>
- [13] Hongyang Gao, Yongjun Chen, and Shuiwang Ji. 2019. Learning Graph Pooling and Hybrid Convolutional Operations for Text Representations. *preprint abs/1901.06965* (2019). arXiv:1901.06965 <http://arxiv.org/abs/1901.06965>
- [14] Hongyang Gao and Shuiwang Ji. 2019. Graph U-Net. *ICLR 2019 Conference Blind Submission* (2019). <https://openreview.net/forum?id=HJePProAct7>
- [15] Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. 2017. Neural Message Passing for Quantum Chemistry. In *ICML*. arXiv:1704.01212. <http://arxiv.org/abs/1704.01212v2>
- [16] G.B. Goh, C. Siegel, A. Vishnu, N. Hodas, and N. Baker. 2018. How Much Chemistry Does a Deep Neural Network Need to Know to Make Accurate Predictions?. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*. 1340–1349. <https://doi.org/10.1109/WACV.2018.00151>
- [17] Johannes Hachmann, Roberto Olivares-Amaya, Sule Atahan-Evrenk, Carlos Amador-Bedolla, Roel S. Sánchez-Carrera, Aryeh Gold-Parker, Leslie Vogt, Anna M. Brockway, and Alán Aspuru-Guzik. 2011. The Harvard Clean Energy Project: Large-Scale Computational Screening and Design of Organic Photovoltaics on the World Community Grid. *The Journal of Physical Chemistry Letters* 2, 17 (aug 2011), 2241–2251. <https://doi.org/10.1021/jz200866s>
- [18] Sabrina Jaeger, Simone Fulle, and Samo Turk. 2018. Mol2vec: Unsupervised Machine Learning Approach with Chemical Intuition. *Journal of Chemical Information and Modeling* 58, 1 (2018), 27–35. <https://doi.org/10.1021/acs.jcim.7b00616> arXiv:https://doi.org/10.1021/acs.jcim.7b00616 PMID: 29268609.
- [19] Junying Li, Deng Cai, and Xiaofei He. 2017. Learning Graph-Level Representation for Drug Discovery. (2017). arXiv:1709.03741 <http://arxiv.org/abs/1709.03741>
- [20] Junying Li, Deng Cai, and Xiaofei He. 2017. Learning Graph-Level Representation for Drug Discovery. *CoRR abs/1709.03741* (2017). arXiv:1709.03741 <http://arxiv.org/abs/1709.03741>
- [21] Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Amey Talwalkar. 2018. Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization. *Journal of Machine Learning Research* 18, 185 (2018), 1–52. <http://jmlr.org/papers/v18/li16-558.html>
- [22] Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard S. Zemel. 2016. Gated Graph Sequence Neural Networks. *International Conference on Learning Representations (ICLR)* (2016). arXiv:1511.05493 <http://arxiv.org/abs/1511.05493>
- [23] Henry W. Lin, Max Tegmark, and David Rolnick. 2017. Why Does Deep and Cheap Learning Work So Well? *Journal of Statistical Physics* 168, 6 (jul 2017), 1223–1247. <https://doi.org/10.1007/s10955-017-1836-5>
- [24] Steven A. Lopez, Benjamin Sanchez-Lengeling, Julio de Goes Soares, and Alán Aspuru-Guzik. 2017. Design Principles and Top Non-Fullerene Acceptor Candidates for Organic Photovoltaics. *Joule* 1, 4 (dec 2017), 857–870. <https://doi.org/10.1016/j.joule.2017.10.006>
- [25] Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodolà, Jan Svoboda, and Michael M. Bronstein. 2017. Geometric deep learning on graphs and manifolds using mixture model CNNs. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017). arXiv:1611.08402 <http://arxiv.org/abs/1611.08402>
- [26] Damian Mrowca, Chengxu Zhuang, Elias Wang, Nick Haber, Li Fei-Fei, Joshua B. Tenenbaum, and Daniel L. K. Yamins. 2018. Flexible Neural Representation for Physics Prediction. *Advances in Neural Information Processing Systems (NIPS)* 31 (2018). arXiv:1806.08047 <http://arxiv.org/abs/1806.08047>
- [27] Mathias Niepert, Mohamed Ahmed, and Konstantin Kutzkov. 2016. Learning Convolutional Neural Networks for Graphs. *Proceedings of the International Conference on Machine Learning (ICML)* 48 (2016). arXiv:1605.05273 <http://arxiv.org/abs/1605.05273>
- [28] M. C. Scharber, D. Mühlbacher, M. Koppe, P. Denk, C. Waldauf, A. J. Heeger, and C. J. Brabec. 2006. Design Rules for Donors in Bulk-Heterojunction Solar Cells—Towards 10 % Energy-Conversion Efficiency. *Advanced Materials* 18, 6 (mar 2006), 789–794. <https://doi.org/10.1002/adma.200501717>
- [29] Chao Shang, Qingqing Liu, Ko-Shin Chen, Jiangwen Sun, Jin Lu, Jinfeng Yi, and Jinbo Bi. 2018. Edge Attention-based Multi-Relational Graph Convolutional Networks. *arXiv e-prints*, Article arXiv:1802.04944 (Feb 2018), arXiv:1802.04944 pages. arXiv:stat.ML/1802.04944 <https://arxiv.org/pdf/1802.04944v1.pdf>
- [30] Martin Simonovsky and Nikos Komodakis. 2017. Dynamic Edge-Conditioned Filters in Convolutional Neural Networks on Graphs. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017), 29–38. <http://arxiv.org/abs/1704.02901v3>
- [31] Gregor Urban, Niranjana Subrahmanya, and Pierre Baldi. 2018. Inner and Outer Recursive Neural Networks for Chemoinformatics Applications. *Journal of Chemical Information and Modeling* 58, 2 (2018), 207–211. <https://doi.org/10.1021/acs.jcim.7b00384> arXiv:https://doi.org/10.1021/acs.jcim.7b00384 PMID: 29320180.
- [32] Robin Winter, Floriane Montanari, Frank NoÁl, and Djork-ArnÁl Clevert. 2019. Learning continuous and data-driven molecular descriptors by translating equivalent chemical representations. *Chem. Sci.* 10 (2019), 1692–1701. Issue 6. <https://doi.org/10.1039/C8SC04175J>
- [33] Zhenqin Wu, Bharath Ramsundar, Evan N. Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S. Pappu, Karl Leswing, and Vijay S. Pande. 2018. MoleculeNet: A Benchmark for Molecular Machine Learning. *Chemical Science* 2 (2018). arXiv:1703.00564 <http://arxiv.org/abs/1703.00564>
- [34] Rex Ying, Jiaxuan You, Christopher Morris, Xiang Ren, William L. Hamilton, and Jure Leskovec. 2018. Hierarchical Graph Representation Learning with Differentiable Pooling. *Advances in Neural Information Processing Systems (NIPS)* 31 (2018). arXiv:1806.08804 <http://arxiv.org/abs/1806.08804>
- [35] Shuangjia Zheng, Xin Yan, Yuedong Yang, and Jun Xu. 2019. Identifying Structure-Property Relationships through SMILES Syntax Analysis with Self-Attention Mechanism. *Journal of Chemical Information and Modeling* 59, 2 (2019), 914–923. <https://doi.org/10.1021/acs.jcim.8b00803> arXiv:https://doi.org/10.1021/acs.jcim.8b00803 PMID: 30669836.

A SUPPLEMENTARY MATERIAL

A.1 Material science application: Clean Energy Project 2017 dataset

In this section, we propose a regression benchmark for hierarchical models using the 2017 non-fullerene electron-acceptor update [24] to the Clean Energy Project molecular library [17]. We refer to this dataset as CEP-2017. This dataset was generated by combining molecular fragments from a reference library generating 51256 unique molecules. These molecular graphs were then used as input to density functional theory electronic-structure calculations of quantum-mechanical observables (such as GAP and HOMO). Restrictions of the crowd-sourced computing platform limited these structures to molecules of 306 electrons or less. The direct observables quantities are then used in a physically motivated but empirical Scharber [28] model to predict power conversion efficiency (PCE). This efficiency is the ultimate figure of merit for a new photovoltaic material.

We emphasize that this data, generated with an approximate density functional theory method, and then used in an empirical PCE model, lacks predictive power in terms of design of new materials. However a machine learning model built on this data is likely to be transferable to other molecular datasets built on higher level theory (such as coupled-cluster calculations) or experimental ground truth. As we are anticipating this future application of the method, we use the raw (`_calc`) values rather than the Gaussian process regressed (to a small experimental dataset) values (`_calib`).

The method of construction of the dataset allows us to highlight the coarse-graining interpretation of the pooling layers introduced in the main text, in terms of the explicit combinatorial building blocks of the non-fullerene electron acceptors.

In Table 3, we show multi-task and single-task test set evaluation R^2 results for the power conversion efficiency (PCE), the band gap (GAP), and the highest occupied molecular orbital (HOMO) energy. We used a dual-message graph-convolutional model with three graph-convolutional layers with node channel dimensions [512, 512, 512] and edge channel dimensions [128, 128, 128] with two interleaved layers of simple pooling. We found our model to be a powerful predictor of both fundamental quantum-mechanical properties (GAP and HOMO), and to a lesser extend the more empirical PCE figure. The inclusion of pooling layers resulted in a significant speedup and only a very mild decay in performance.

pooling ratio	Multi-task			Single-task
	R^2 on PCE	R^2 on GAP	R^2 on HOMO	R^2 on PCE
none	0.862 ± 0.005	0.967 ± 0.001	0.981 ± 0.000	0.866 ± 0.003
0.9	0.863 ± 0.003	0.966 ± 0.001	0.981 ± 0.000	0.862 ± 0.002
0.8	0.860 ± 0.003	0.966 ± 0.001	0.981 ± 0.001	0.859 ± 0.004
0.7	0.856 ± 0.003	0.964 ± 0.001	0.980 ± 0.001	0.855 ± 0.004
0.6	0.854 ± 0.007	0.962 ± 0.002	0.979 ± 0.001	0.853 ± 0.003
0.5	0.844 ± 0.003	0.955 ± 0.002	0.974 ± 0.001	0.833 ± 0.007
	RMSE on PCE	RMSE on GAP	RMSE on HOMO	RMSE on PCE
none	0.217 ± 0.004	0.177 ± 0.002	0.134 ± 0.001	0.215 ± 0.002
0.9	0.217 ± 0.002	0.180 ± 0.002	0.134 ± 0.001	0.218 ± 0.002
0.8	0.220 ± 0.002	0.179 ± 0.002	0.135 ± 0.003	0.220 ± 0.003
0.7	0.179 ± 0.097	0.148 ± 0.080	0.112 ± 0.060	0.223 ± 0.003
0.6	0.224 ± 0.005	0.190 ± 0.005	0.143 ± 0.004	0.225 ± 0.003
0.5	0.232 ± 0.002	0.208 ± 0.004	0.159 ± 0.003	0.239 ± 0.005

Table 3: Multi-task and single-task benchmark R^2 results for power conversion efficiency (PCE), band gap (GAP), and highest occupied molecular orbital (HOMO) energy of the CEP-2017 benchmark for different ratios of kept nodes in each pooling step (averaged over 5 runs, with 5 random splits). Speedup of pooling runs is measured in terms of elapsed real time compared to the run without pooling.

A.2 Pooling layers illustrations

In Fig. 2(a-c) we visualize the effect of two consecutive pooling layers (each keeping only 50% of the nodes) on a batch of molecules for a DM-SIMPLEPOOLING model trained on a random split of the CEP-2017 dataset introduced in Sec. A.1. After the first pooling layer (Fig. 2(b)), the model has approximately learned to group rings and identify the backbones or main connected chains of the molecules. After the second pooling layer (Fig. 2(c)), the molecular graphs have been reduced to basic, abstract components connected by chains, encoding a coarse-grained representation of the original molecules. Disconnected parts can be interpreted as a consequence of the aggressive pooling forcing the model to pay attention to the parts it considers most relevant for the task at hand.

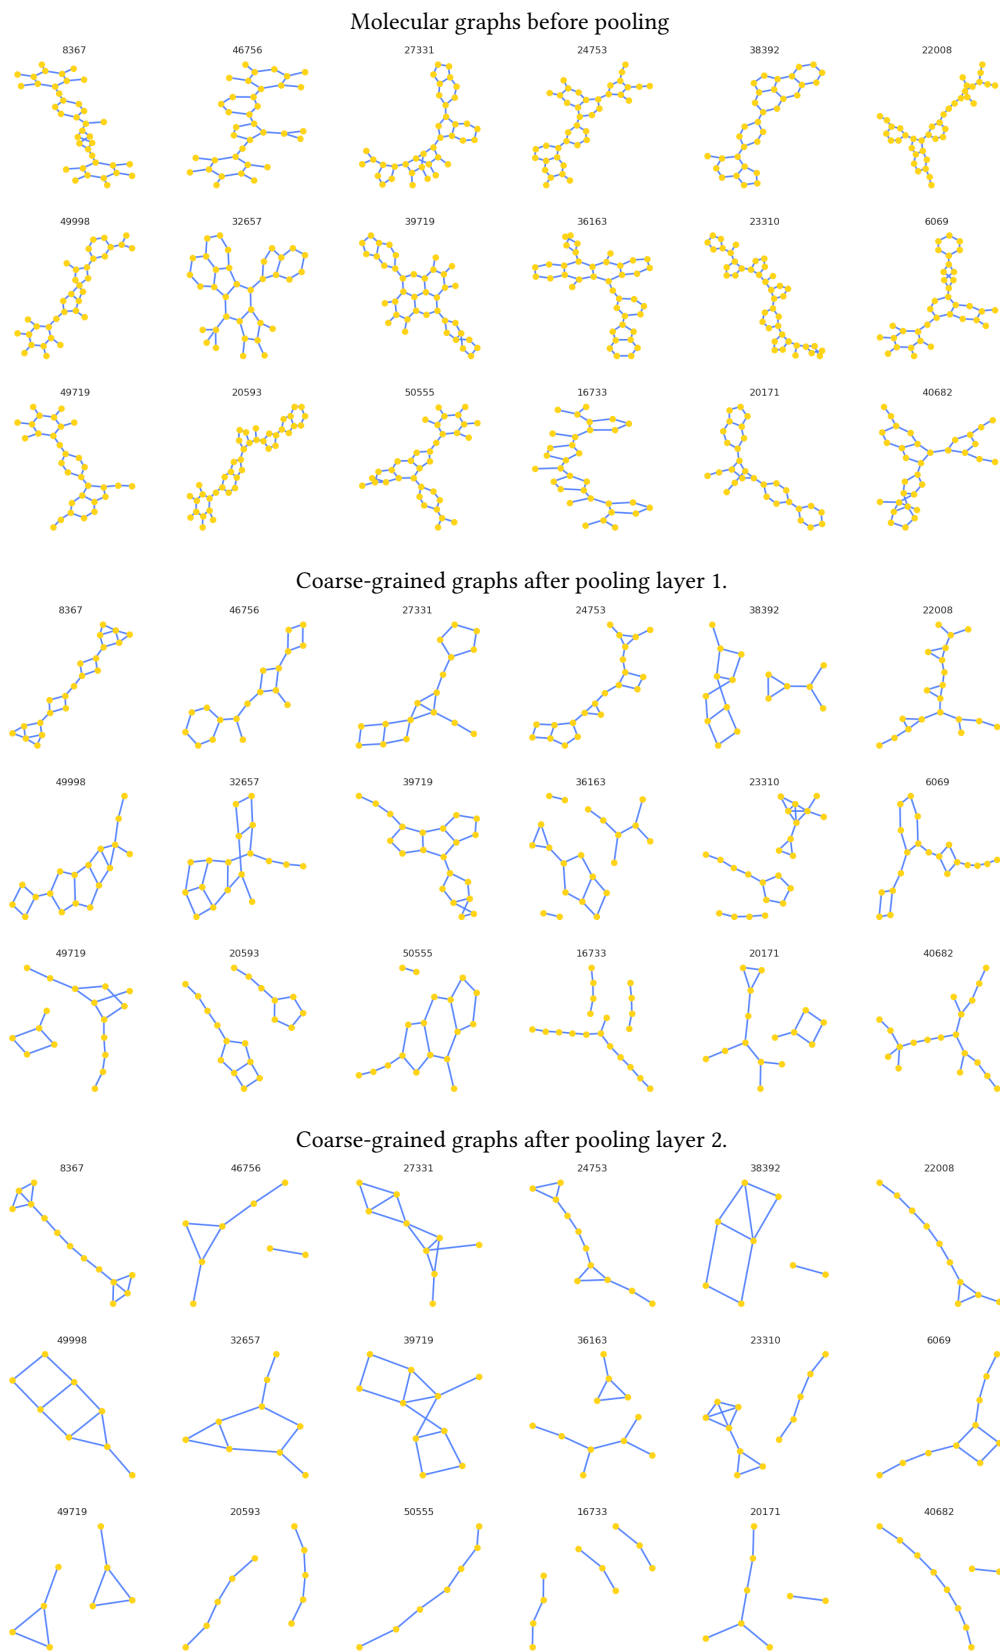


Figure 2: Pooling of molecular graphs (heavy atoms only) sampled from CEP-2017 dataset.

Dataset	Model	Keep ratio	Node Channels	Edge Channels	
<i>ESOL</i>	NoPooling		[128, 128]	[128, 128]	
	CoarseGrainPooling	0.9	[128, 128]	[256, 256]	
		0.8	[128, 128]	[64, 64]	
		0.7	[512, 512, 512]	[128, 128, 128]	
		0.6	[256, 256, 256]	[128, 128, 128]	
		0.5	[128, 128]	[64, 64]	
	SimplePooling	0.9	[256, 256]	[128, 128]	
		0.8	[256, 256, 256]	[256, 256, 256]	
		0.7	[128, 128, 128]	[128, 128, 128]	
		0.6	[512, 512]	[128, 128]	
		0.5	[256, 256]	[256, 256]	
	<i>Lipophilicity</i>	NoPooling		[256, 256, 256]	[64, 64, 64]
		CoarseGrainPooling	0.9	[256, 256]	[128, 128]
			0.8	[256, 256]	[64, 64]
			0.7	[256, 256]	[64, 64]
0.6			[256, 256]	[256, 256]	
0.5			[512, 512]	[64, 64]	
SimplePooling		0.9	[512, 512]	[64, 64]	
		0.8	[128, 128]	[128, 128]	
		0.7	[256, 256]	[128, 128]	
		0.6	[512, 512]	[128, 128]	
		0.5	[256, 256]	[128, 128]	
<i>BBBP</i>		NoPooling		[128, 128]	[256, 256]
		CoarseGrainPooling	0.9	[256, 256]	[256, 256]
			0.8	[512, 512]	[256, 256]
			0.7	[128, 128, 128]	[128, 128, 128]
	0.6		[256, 256]	[64, 64]	
	0.5		[256, 256, 256]	[64, 64, 64]	
	SimplePooling	0.9	[128, 128, 128]	[256, 256, 256]	
		0.8	[512, 512]	[128, 128]	
		0.7	[256, 256]	[64, 64]	
		0.6	[512, 512]	[256, 256]	
		0.5	[128, 128, 128]	[64, 64, 64]	
	<i>HIV</i>	All models		[512, 512, 512]	[128, 128, 128]

Table 4: Graph-convolutional model hyperparameters used in this work.