
Towards explainable message passing networks for predicting carbon dioxide adsorption in metal-organic frameworks

Ali Raza^{*1}, Faaiq Waqar¹, Arni Sturluson², Cory Simon^{†2} and Xiaoli Fern^{‡1}

¹School of Electrical Engineering and Computer Science, Oregon State University

²School of Chemical, Biological, and Environmental Engineering, Oregon State University

Abstract

Metal-organic framework (MOFs) are nanoporous materials that could be used to capture carbon dioxide from the exhaust gas of fossil fuel power plants to mitigate climate change. In this work, we design and train a message passing neural network (MPNN) to predict simulated CO₂ adsorption in MOFs. Towards providing insights into what substructures of the MOFs are important for the prediction, we introduce a soft attention mechanism into the readout function that quantifies the contributions of the node representations towards the graph representations. We investigate different mechanisms for sparse attention to ensure only the most relevant substructures are identified.

1 Introduction

Anthropogenic carbon dioxide (CO₂) emissions are a major contributor to climate change and ocean acidification [1]. Carbon dioxide capture and storage [2] is among a concerted portfolio of approaches [3] to stabilize and eventually reduce our CO₂ emissions. In post-combustion carbon capture, CO₂ is separated from the combustion exhaust gas of fossil fuel power plants, at the point of emission, and then geologically sequestered [1]. Metal-organic frameworks (MOFs) [4] are nano-porous, crystalline materials that can selectively adsorb CO₂ [5, 6] and therefore could be used to capture CO₂ from the flue gas of fossil fuel power plants [7].

MOFs are acclaimed as “designer materials” [8] because the chemistry of the internal surface of the MOF can be (computationally) designed to target the adsorption of CO₂ [9]. MOFs are synthesized modularly, by linking organic molecules to metals/metal clusters to form an extended network. Due to the abundance of molecular building blocks and their post-synthetic modifiability, the space of MOFs is vast. Molecular models and simulations [10, 11] and machine learning [12, 13, 14] play an important role in navigating this vast space of MOFs to find a suitable/optimal MOF for energy-efficient CO₂ capture and release [15].

Here, we design and train a message passing neural network (MPNN) [16, 17] to predict the (simulated) amount of CO₂ adsorption in MOFs. As opposed to the traditional machine learning approach of human-engineering a feature vector to represent the structure of the MOF [18, 19, 20, 21, 22], the MPNN directly takes a graph representation of the MOF structure as input and automatically learns a vector representation of the MOF to use for the prediction task, in an end-to-end manner. This is achieved by iteratively passing information between neighboring nodes to learn hidden

*razaa@oregonstate.edu

†cory.simon@oregonstate.edu

‡xfern@oregonstate.edu

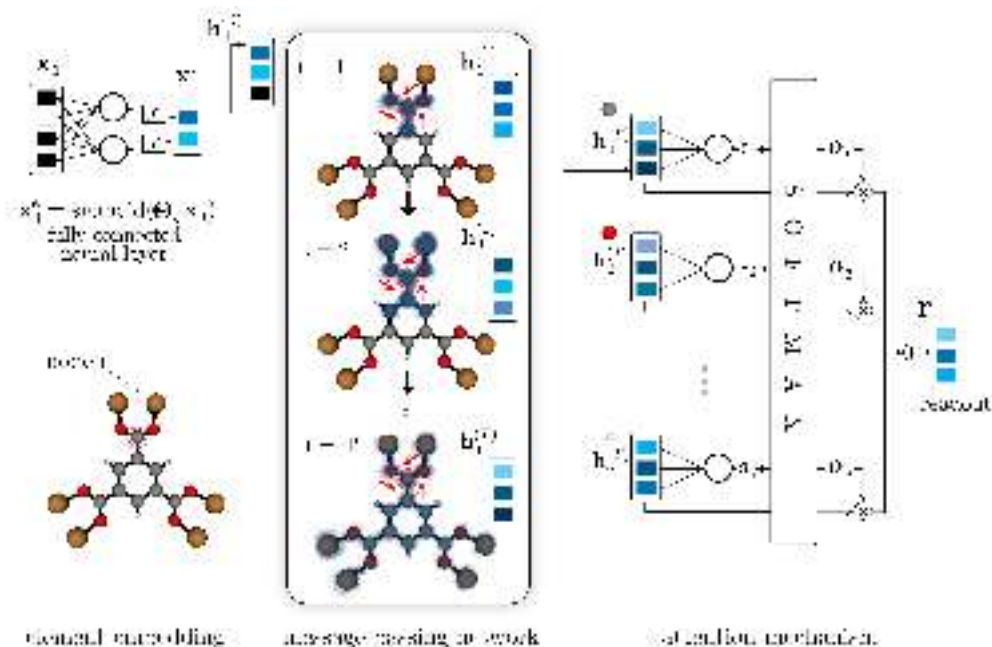


Figure 1: The architecture of our message passing neural network (MPNN).

representations of the local bonding environments within the MOFs, then, through a readout function, aggregating the local node representations into a graph representation used for the prediction task. In a step towards explaining the predictions of the MPNN by identifying important substructures in the graph, we incorporated an attention mechanism in the readout function of the MPNN that quantifies the contribution of each node’s representation to the graph representation. Explainability is advantageous because it (i) can elucidate design rules and chemical intuition for synthesizing MOFs with desirable adsorption properties and (ii) build appropriate trust/skepticism of particular predictions based on the explanation.

MPNNs [16, 17] have been used to predict the properties of molecules and materials [23, 24, 25, 17, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35], as well as to generate molecules and materials with desired properties [36, 37]. There has been limited efforts in interpreting/explaining MPNNs or graph neural networks (GNNs) in general [29, 38, 39, 40, 41, 42].

2 Proposed framework

2.1 Problem overview

We aim to predict the equilibrium CO_2 adsorption in a MOF at a given temperature and pressure, $a \in \mathbb{R}^+$ [mmol/g]. Each MOF structure is represented as an undirected, node-labeled graph $G = (\mathcal{V}, \mathcal{E}, \mathbf{X})$, where \mathcal{V} is the set of $n = |\mathcal{V}|$ nodes or vertices, representing atoms, \mathcal{E} is the set of edges, representing bonds, and $\mathbf{X} \in \mathbb{R}^{d \times n}$ is the node feature matrix, whose columns are one-hot encodings of the chemical elements of the atoms (d possible elements). In a supervised manner, we aim to learn a function that maps a MOF to its predicted CO_2 adsorption: $f : G \mapsto f(G) = a$.

2.2 Converting a MOF crystal structure to a graph

We constructed the node-labeled graph G representing each MOF from its unit cell, the list of atoms in the unit cell, and the crystallographic coordinates of those atoms. A bijection exists between the nodes \mathcal{V} and the atoms comprising the unit cell of the MOF. Two atoms are assigned an edge (bond) iff, as in Refs. [34, 43, 44], (i) they are less than a distance r apart, with r the sum of their covalent radii [45] (some metals modified) plus a 0.25 \AA tolerance, and (ii) they share a Voronoi face in a

Voronoi diagram of the surrounding atoms. We used periodic distance in our calculations to include bonds across the periodic boundary of the unit cell.

2.3 Message passing neural network that represents f

Fig. 1 shows the architecture of our message passing neural network (MPNN) that includes an attention mechanism to construct the graph representation from the set of node representations.

Message passing. First, our MPNN operates on nodes and learns a vector representation of the local bonding environment of each node in the graph. This is achieved by a chemical element embedding followed by iterations of message passing.

The chemical element embedding layer maps the one-hot encoding of the chemical element of the node to a low-dimensional, dense representation:

$$\mathbf{x}_v^e = \text{sigmoid}(\Theta_e \mathbf{x}_v), \quad (1)$$

where \mathbf{x}_v is column v of \mathbf{X} and Θ_e is a $r \times d$ learned matrix, with $r < d$.

Message passing is then used to learn a representation of each node encapsulating information about its local bonding environment. Let $\mathbf{h}_v^{(t)} \in \mathbb{R}^k$ be the hidden representation of node v at time step t , with $k \geq r$, initialized with the element embedding \mathbf{x}_v^e padded with zeros. In each time step, every node receives information from its neighbors and updates its hidden representation accordingly, from an aggregated message, using a gated graph neural network (GGNN) [46]. The aggregated message received by node v is:

$$\mathbf{m}_v^{(t+1)} = \Theta_m \sum_{u \in \mathcal{N}(v)} \mathbf{h}_u^{(t)} \quad (2)$$

where Θ_m is a learned $k \times k$ matrix shared across all nodes and $\mathcal{N}(v)$ is the set of nodes that share an edge with node v . A Gated Recurrent Unit (GRU) (shared across all nodes) then updates the hidden representation of node v :

$$\mathbf{h}_v^{(t+1)} = \text{GRU}(\mathbf{h}_v^{(t)}, \mathbf{m}_v^{(t+1)}). \quad (3)$$

We conduct T time steps of message passing, after which $\mathbf{h}_v^{(T)}$ contains information about the bonding environment of node v within a graph-distance of T .

Readout and prediction. A readout function [47] maps the set of hidden node representations to a fixed-size vector representation of the entire graph. We use a soft attention mechanism [46] where the attention of node v , α_v , is computed from the hidden features via a `softmax`:

$$\alpha_v = \frac{e^{\theta_s^T \mathbf{h}_v^{(T)}}}{\sum_{u=1}^n e^{\theta_s^T \mathbf{h}_u^{(T)}}} \quad (4)$$

where $\theta_s \in \mathbb{R}^k$ is a learned vector shared across all nodes. To introduce sparsity, we also try (1) the quasi-norm $L(0.5)$ regularization of the attention scores and (2) `sparsemax` [48] in place of `softmax`. The attention score of node v then determines the contribution of its hidden representation $\mathbf{h}_v^{(T)}$ to the representation of the graph, \mathbf{r} :

$$\mathbf{r} = \sum_{v=1}^n \alpha_v \mathbf{h}_v^{(T)}. \quad (5)$$

Finally, a neural network predicts the CO_2 adsorption (\hat{a}) from the graph representation:

$$\hat{a} = \text{softplus}(\theta_a^T (\text{sigmoid}(\Theta_a \mathbf{r}))) \quad (6)$$

where Θ_a is a learned $z \times k$ matrix and $\theta_a \in \mathbb{R}^z$ is a learned vector. The `softplus` ensures $a > 0$.

3 Results

As train, test, and validation data, we use simulated CO_2 uptake at 298 K and 0.15 bar from Ref. [49], taken from the Materials Cloud [50], in 6 103 computation-ready, experimental MOF structures [51].

Method	mean (std)			
	MAD	MSE	ρ_r	Entropy (sparsity)
MPNN (softmax)	0.616 (0.03)	0.868 (0.10)	0.764 (0.02)	0.78 (0%)
MPNN (sparsemax)	0.666 (0.04)	1.000 (0.11)	0.732 (0.02)	0.45 (94%)
MPNN (L0.5, $\lambda = 0.001$)	0.645 (0.02)	0.933 (0.08)	0.743 (0.01)	0.72 (0%)
MPNN (L0.5, $\lambda = 0.05$)	0.737 (0.04)	0.1.174 (0.13)	0.684 (0.04)	0.32 (0%)

Table 1: Prediction performance and attention sparsity by different methods. Mean and standard deviation (std) over 10 folds.

We use the mean absolute deviation (MAD) loss function $\ell = \frac{1}{M} \sum_{m=1}^M \|\hat{a}_m - a_m\|_1$ to train our network within $K = 10$ -fold cross validation, where M is the total number of MOFs, \hat{a}_m is the predicted CO₂ adsorption of MOF m predicted by the MPNN by eqn. 5, a_m is the simulated CO₂ adsorption (treated as ground truth) and $\|\cdot\|_1$ is the L1 norm. Through hyperparameter exploration, we settled on $r = 10$, $k = 70$, and $T = 4$. Tab. 1 summarizes the performance of our model using the mean absolute deviation (MAD), mean square error (MSE), Spearman’s rank correlation coefficient, ρ_r , and normalized entropy of the attentions (1 for uniform attention across all nodes and 0 for all attention concentrating on one node). Vanilla softmax is able to achieve the best MAD performance. Fig. 2a shows a parity plot, using softmax, for the test MOFs during the cross-validation procedure. Sparsemax introduced substantial sparsity in the attention scores (%94 of the attention scores are zero); however, there is no way to control the sparsity. Using Quasi-norm L(0.5) regularisation produces a less uniform attention distribution, but with a price of slightly higher MAD. The regularization parameter, λ , enables us to trade-off training set accuracy with sparsity. Smaller entropy ($\lambda = 0.001$ to 0.05) results in less accuracy (MAD = 0.645 to 0.737).

4 Discussion: towards explainability

We include the attention score α_v as a step towards an MPNN with explainable predictions. If α_v is large, the hidden representation of node v had a significant contribution to the graph representation \mathbf{r} used to predict adsorption, a . Fig. 2c visualizes the attention of each node in a MOF as an example. The local bonding environments of the darker atoms contributed more to the final graph representation used for the prediction task than the lighter atoms.




Figure 2: (a) Parity plot for the MPNN (softmax), showing the MPNN-predicted vs. simulated CO₂ adsorption, including every MOF in the data set when it served as a test MOF in cross validation. Diagonal line shows equality. (b, c) Visualization of MOF IRMOF-1, where color indicates (b) the identity of the atom and (c) the attention score α_v on the atom using the MPNN (softmax). Black boxes = unit cell.

Our inspection of the attention scores across different MOFs did not yield any clear insight or chemically meaningful patterns. This leads us to believe that the attention score is not as meaningful as we have hoped for identifying important substructures. we further believe this is due to (i) $\mathbf{h}_v^{(T)}$ of node v is enriched with information from all nodes within T steps from node v via message

passing; (ii) neighboring nodes tend to have strong similarity; and (iii) the attention scores do not necessarily reflect the importance of node v for the predicted adsorption [52]. Therefore, we are currently working to redesign the MPNN to short-circuit message passing, isolate the individual contributions of the nodes, and rigorously lend more explainability of the prediction.

Acknowledgments and Disclosure of Funding

The authors acknowledge the National Science Foundation for support under grants No. 1920945 and No. 1521687.

References

- [1] Berend Smit, Jeffrey A Reimer, Curtis M Oldenburg, and Ian C Bourg. *Introduction to carbon capture and sequestration*, volume 1. World Scientific, 2014.
- [2] R Stuart Haszeldine. Carbon capture and storage: how green can black be? *Science*, 325(5948):1647–1652, 2009.
- [3] Stephen Pacala and Robert Socolow. Stabilization wedges: solving the climate problem for the next 50 years with current technologies. *Science*, 305(5686):968–972, 2004.
- [4] Hiroyasu Furukawa, Kyle E Cordova, Michael O’Keeffe, and Omar M Yaghi. The chemistry and applications of metal-organic frameworks. *Science*, 341(6149):1230444, 2013.
- [5] Kenji Sumida, David L Rogow, Jarad A Mason, Thomas M McDonald, Eric D Bloch, Zoey R Herm, Tae-Hyun Bae, and Jeffrey R Long. Carbon dioxide capture in metal-organic frameworks. *Chemical Reviews*, 112(2):724–781, 2012.
- [6] Deanna M D’Alessandro, Berend Smit, and Jeffrey R Long. Carbon dioxide capture: prospects for new materials. *Angewandte Chemie International Edition*, 49(35):6058–6082, 2010.
- [7] Alexander Schoedel, Zhe Ji, and Omar M Yaghi. The role of metal-organic frameworks in a carbon-neutral energy cycle. *Nature Energy*, 1(4):1–13, 2016.
- [8] Christopher H Hendon, Adam J Rieth, Maciej D Korzyński, and Mircea Dincă. Grand challenges and future opportunities for metal-organic frameworks. *ACS Central Science*, 3(6):554–563, 2017.
- [9] Peter G Boyd, Arunraj Chidambaram, Enrique García-Díez, Christopher P Ireland, Thomas D Daff, Richard Bounds, Andrzej Gładysiak, Pascal Schouwink, Seyed Mohamad Moosavi, M Mercedes Maroto-Valer, et al. Data-driven design of metal-organic frameworks for wet flue gas co₂ capture. *Nature*, 576(7786):253–256, 2019.
- [10] Arni Sturluson, Melanie T Huynh, Alec R Kaija, Caleb Laird, Sunghyun Yoon, Feier Hou, Zhenxing Feng, Christopher E Wilmer, Yamil J Colón, Yongchul G Chung, Siderius D, and Simon C. The role of molecular modelling and simulation in the discovery and deployment of metal-organic frameworks for gas storage and separation. *Molecular Simulation*, 45(14-15):1082–1121, 2019.
- [11] Peter G Boyd, Yongjin Lee, and Berend Smit. Computational development of the nanoporous materials genome. *Nature Reviews Materials*, 2(8):1–15, 2017.
- [12] Sanggyu Chong, Sangwon Lee, Baekjun Kim, and Jihan Kim. Applications of machine learning in metal-organic frameworks. *Coordination Chemistry Reviews*, 423:213487, 2020.
- [13] Kevin Maik Jablonka, Daniele Ongari, Seyed Mohamad Moosavi, and Berend Smit. Big-data science in porous materials: Materials genomics and machine learning. *Chemical Reviews*, 2020.
- [14] Zenan Shi, Wenyuan Yang, Xiaomei Deng, Chengzhi Cai, Yaling Yan, Hong Liang, Zili Liu, and Zhiwei Qiao. Machine-learning-assisted high-throughput computational screening of high performance metal-organic frameworks. *Molecular Systems Design & Engineering*, 5(4):725–742, 2020.

- [15] Johanna M Huck, Li-Chiang Lin, Adam H Berger, Mahdi Niknam Shahrak, Richard L Martin, Abhoyjit S Bhowan, Maciej Haranczyk, Karsten Reuter, and Berend Smit. Evaluating different classes of porous materials for carbon capture. *Energy & Environmental Science*, 7(12):4132–4146, 2014.
- [16] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S Yu. A comprehensive survey on graph neural networks. *arXiv preprint arXiv:1901.00596*, 2019.
- [17] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1263–1272. JMLR. org, 2017.
- [18] Cory M Simon, Rocio Mercado, Sondre K Schnell, Berend Smit, and Maciej Haranczyk. What are the best materials to separate a xenon/krypton mixture? *Chemistry of Materials*, 27(12):4459–4475, 2015.
- [19] Maryam Pardakhti, Ehsan Moharreri, David Wanik, Steven L Suib, and Ranjan Srivastava. Machine learning using combined structural and chemical descriptors for prediction of methane adsorption performance of metal organic frameworks (mofs). *ACS Combinatorial Science*, 19(10):640–645, 2017.
- [20] Ryther Anderson, Jacob Rodgers, Edwin Argueta, Achay Biong, and Diego A Gómez-Gualdrón. Role of pore chemistry and topology in the co₂ capture capabilities of mofs: from molecular simulation to machine learning. *Chemistry of Materials*, 30(18):6325–6337, 2018.
- [21] Hana Dureckova, Mykhaylo Krykunov, Mohammad Zein Aghaji, and Tom K Woo. Robust machine learning models for predicting high co₂ working capacity and co₂/h₂ selectivity of gas adsorption in metal organic frameworks for precombustion carbon capture. *The Journal of Physical Chemistry C*, 123(7):4133–4139, 2019.
- [22] Benjamin J Bucior, N Scott Bobbitt, Timur Islamoglu, Subhadip Goswami, Arun Gopalan, Taner Yildirim, Omar K Farha, Neda Bagheri, and Randall Q Snurr. Energy-based descriptors to rapidly predict hydrogen storage in metal–organic frameworks. *Molecular Systems Design & Engineering*, 4(1):162–174, 2019.
- [23] Kevin Yang, Kyle Swanson, Wengong Jin, Connor Coley, Philipp Eiden, Hua Gao, Angel Guzman-Perez, Timothy Hopper, Brian Kelley, Miriam Mathea, et al. Analyzing learned molecular representations for property prediction. *Journal of Chemical Information and Modeling*, 59(8):3370–3388, 2019.
- [24] Steven Kearnes, Kevin McCloskey, Marc Berndl, Vijay Pande, and Patrick Riley. Molecular graph convolutions: moving beyond fingerprints. *Journal of Computer-aided Molecular Design*, 30(8):595–608, 2016.
- [25] Jonathan M Stokes, Kevin Yang, Kyle Swanson, Wengong Jin, Andres Cubillos-Ruiz, Nina M Donghia, Craig R MacNair, Shawn French, Lindsey A Carfrae, Zohar Bloom-Ackerman, et al. A deep learning approach to antibiotic discovery. *Cell*, 180(4):688–702, 2020.
- [26] David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P Adams. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in Neural Information Processing Systems*, pages 2224–2232, 2015.
- [27] Connor W Coley, Regina Barzilay, William H Green, Tommi S Jaakkola, and Klavs F Jensen. Convolutional embedding of attributed molecular graphs for physical property prediction. *Journal of Chemical Information and Modeling*, 57(8):1757–1772, 2017.
- [28] Peter Bjørn Jørgensen, Karsten Wedel Jacobsen, and Mikkel N Schmidt. Neural message passing with edge updates for predicting properties of molecules and materials. *arXiv preprint arXiv:1806.03146*, 2018.
- [29] Bowen Tang, Skyler T Kramer, Meijuan Fang, Yingkun Qiu, Zhen Wu, and Dong Xu. A self-attention based message passing neural network for predicting molecular lipophilicity and aqueous solubility. *Journal of Cheminformatics*, 12(1):1–9, 2020.

- [30] Benjamin Sanchez-Lengeling, Jennifer N Wei, Brian K Lee, Richard C Gerkin, Alán Aspuru-Guzik, and Alexander B Wiltschko. Machine learning for scent: Learning generalizable perceptual representations of small molecules. *arXiv preprint arXiv:1910.10685*, 2019.
- [31] Chi Chen, Weike Ye, Yunxing Zuo, Chen Zheng, and Shyue Ping Ong. Graph networks as a universal machine learning framework for molecules and crystals. *Chemistry of Materials*, 31(9):3564–3572, 2019.
- [32] Peter C St. John, Caleb Phillips, Travis W Kemper, A Nolan Wilson, Yanfei Guan, Michael F Crowley, Mark R Nimlos, and Ross E Larsen. Message-passing neural networks for high-throughput polymer screening. *The Journal of Chemical Physics*, 150(23):234111, 2019.
- [33] Tian Xie and Jeffrey C Grossman. Hierarchical visualization of materials space with graph convolutional neural networks. *The Journal of Chemical Physics*, 149(17):174111, 2018.
- [34] Tian Xie and Jeffrey C Grossman. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Physical Review Letters*, 120(14):145301, 2018.
- [35] Cheol Woo Park and Chris Wolverton. Developing an improved crystal graph convolutional neural network framework for accelerated materials discovery. *arXiv preprint arXiv:1906.05267*, 2019.
- [36] Rocío Mercado, Tobias Rastemo, Edvard Lindelöf, Günter Klambauer, Ola Engkvist, Hongming Chen, and Esben Jannik Bjerrum. Graph networks for molecular design. *ChemRxiv*, 2020.
- [37] Yujia Li, Oriol Vinyals, Chris Dyer, Razvan Pascanu, and Peter Battaglia. Learning deep generative models of graphs. *arXiv preprint arXiv:1803.03324*, 2018.
- [38] Chris Lin, Gerald J Sun, Krishna C Bulusu, Jonathan R Dry, and Marylens Hernandez. Graph neural networks including sparse interpretability. *arXiv preprint arXiv:2007.00119*, 2020.
- [39] Qiang Huang, Makoto Yamada, Yuan Tian, Dinesh Singh, Dawei Yin, and Yi Chang. Graphlime: Local interpretable model explanations for graph neural networks. *arXiv preprint arXiv:2001.06216*, 2020.
- [40] Chaojie Ji, Ruxin Wang, Ye Li, and Hongyan Wu. Perturb more, trap more: Understanding behaviors of graph neural networks. *arXiv preprint arXiv:2004.09808*, 2020.
- [41] Xue Li and Yuanzhi Cheng. Understanding the message passing in graph neural networks via power iteration. *arXiv preprint arXiv:2006.00144*, 2020.
- [42] Zhitao Ying, Dylan Bourgeois, Jiaxuan You, Marinka Zitnik, and Jure Leskovec. Gnnexplainer: Generating explanations for graph neural networks. In *Advances in neural information processing systems*, pages 9244–9255, 2019.
- [43] Olexandr Isayev, Corey Oses, Cormac Toher, Eric Gossett, Stefano Curtarolo, and Alexander Tropsha. Universal fragment descriptors for predicting properties of inorganic crystals. *Nature Communications*, 8(1):1–12, 2017.
- [44] Ali Raza, Arni Sturluson, Cory M Simon, and Xiaoli Fern. Message passing neural networks for partial charge assignment to metal–organic frameworks. *The Journal of Physical Chemistry C*, 124(35):19070–19082, 2020.
- [45] Beatriz Cordero, Verónica Gómez, Ana E Platero-Prats, Marc Revés, Jorge Echeverría, Eduard Cremades, Flavia Barragán, and Santiago Alvarez. Covalent radii revisited. *Dalton Transactions*, (21):2832–2838, 2008.
- [46] Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493*, 2015.
- [47] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 2020.

- [48] Andre Martins and Ramon Astudillo. From softmax to sparsemax: A sparse model of attention and multi-label classification. In *International Conference on Machine Learning*, pages 1614–1623, 2016.
- [49] Seyed Mohamad Moosavi, Aditya Nandy, Kevin Maik Jablonka, Daniele Ongari, Jon Paul Janet, Peter G Boyd, Yongjin Lee, Berend Smit, and Heather J Kulik. Understanding the diversity of the metal-organic framework ecosystem. *Nature Communications*, 11(1):1–10, 2020.
- [50] Leopold Talirz, Snehal Kumbhar, Elsa Passaro, Aliaksandr V Yakutovich, Valeria Granata, Fernando Gargiulo, Marco Borelli, Martin Uhrin, Sebastiaan P Huber, Spyros Zoupanos, et al. Materials cloud, a platform for open computational science. *arXiv preprint arXiv:2003.12510*, 2020.
- [51] Yongchul G Chung, Emmanuel Haldoupis, Benjamin J Bucior, Maciej Haranczyk, Seulchan Lee, Hongda Zhang, Konstantinos D Vogiatzis, Marija Milisavljevic, Sanliang Ling, Jeffrey S Camp, et al. Advances, updates, and analytics for the computation-ready, experimental metal-organic framework database: Core mof 2019. *Journal of Chemical & Engineering Data*, 64(12):5985–5998, 2019.
- [52] Sarthak Jain and Byron C Wallace. Attention is not explanation. *arXiv preprint arXiv:1902.10186*, 2019.