Applying Transfer Learning to Defect Graph Neural Networks for Defect Formation Energy Predictions

Thomas Bouchard[1], Angela Zhang[2], Justin Garrigus[3], Fatimah Habis[4], Dr. Yuanxi Wang[4]
[1]Department of Physics, Engineering, and Astronomy, Austin Peay State University, Clarksville, TN
[2]College of Natural Sciences, University of Texas at Austin, Austin, TX
[3]Computer Science and Engineering Department, University of North Texas, Denton, TX
[4]Physics Department, University of North Texas, Denton, TX

Predicting defect formation energies in crystal structures is crucial for advancing materials science and has significant implications for quantum computing. This study enhances the predictive capabilities of Crystal Graph Convolutional Neural Networks (CGCNN) by integrating transfer learning techniques. Initially, we compared the standard CGCNN model with a Defect Graph Neural Network (dGNN) specifically adjusted for defect formation energies. We then leveraged a pretrained model on pristine crystal formation energies and applied it to a smaller dataset of defect formation energies. We introduced new pooling functions, including dropout-styled, average, and max pooling, to enhance the model's performance. Additionally, an optimizer was developed to fine-tune command line parameters, ensuring optimal results. Future work aims to explore the potential of an advisor-expert model structure, where multiple sub-models (advisors) specializing in different target quantities contribute to the final prediction of the primary (expert) model. The integration of a supercell representation in our model is also planned, representing the crystal lattice as a large graph that includes both pristine and defect structures. This research has applications in quantum computing, where accurate prediction of defect formation energies can lead to the development of more efficient quantum processors and devices by optimizing the materials used in their construction. These advancements are expected to significantly improve the accuracy and efficiency of defect formation energy predictions, contributing to progress in both materials science and quantum technology.

Keywords: defect formation energies; crystal graph convolutional neural network; pooling functions; transfer learning; defect graph neural network; quantum computing

I. INTRODUCTION

The search for different crystals has always been an important question in the realm of scientific discovery due to their useful properties, and one particular area of importance is the study of defect crystal structures, which contain various impurities such as oxygen vacancies or substitutions. The advancement of such materials is crucial, as they can be used for electrodes [1], as transparent electric conductors [2], and even as qubits [3]. In the past, density functional theory (DFT), aka calculations based on theoretical values of material behavior, has been the most common method used to find such materials, which rely on expensive and slow computations and can take on the scale of months or years to discover 100 new materials [4].

Recent years have seen the rise of machine learning to predict such properties, using specialized graph neural networks. This is usually done by manually setting vectors for materials and picking what one thinks are the most important properties [5], which needs adjusting every time a new property is predicted. Additionally, machine learning models tend to be less accurate than DFT calculations on average for defect material properties, which is likely due to the lack of data regarding materials with defects, and much more data for pristine materials [6]. Transfer learning is a solution: a model can be trained on a large dataset of pristine materials before being fine-tuned on a much smaller database of defect material properties, producing much improved performance.

In this paper, we perform transfer learning on a state-of-the-art defect Graph Neural Network (dGNN) [7] [8]. We do this by pretraining a model on pristine crystal formation energy, and fine-tuning on defect formation energy. We also implemented an optimizer and different cross validation techniques that were not in the original paper, which allowed us to properly verify the accuracy of our model, and showed that leave-two-out was the best optimization method out of the ones we used. We compare the model with last year's transfer learning techniques, which were applied to a different model, and in the future we plan to further refine our techniques to ensure even better results.

II. METHODS

The databases used for this project included a database collected from the materials project for pristine crystal formation energy [9] in which the model was pretrained on, and a database of vacancy formation enthalpies [10] for the model to fine-tune on.

Our primary model, the Defect Graph Neural Network (dGNN), represents each crystal structure as a graph where nodes correspond to atoms and edges to bonds, with nodes characterized by feature vectors containing information such as atomic number, and edges characterized by features such as bond length and bond type. To leverage transfer learning, we pretrained the dGNN on a large dataset of pristine crystal formation energies, using this pretrained model as a foundation for subsequent training on the smaller dataset of defect formation energies.

We experimented with different pooling functions to determine their impact on performance, including baseline dGNN pooling (selecting the parameter vector of the atom to be defected), dropout-styled pooling (randomly selecting a node from the graph and forwarding its vector to the fully connected layers), average pooling (taking the element-wise average of all atom vectors in the graph), and max pooling (taking the element-wise maximum of all atom vectors). Each pooling function was integrated into the dGNN, and their performances were compared.

Additionally, we developed an optimizer to fine-tune command line parameters for running the model. We implemented random optimization, which generates random values corresponding to each parameter, creating a Python command that combines these options, running the Python command, which will run the model, and evaluating the resulting model's performance.

To assess the model's performance, we generated figures illustrating performance improvement over different data sizes, comparison of baseline dGNN results with those of the transfer learning applied dGNN model, and performance comparison of different pooling functions.

III. RESULTS AND DISCUSSIONS

We tested how well transfer learning worked with Defect Graph Neural Networks (dGNNs) to predict defect formation energies in crystal structures. We compared the performance of dGNN models trained from scratch with those enhanced by transfer learning across different dataset sizes (10%, 40%, and 100% of the available data).

Figure 1 shows the Mean Absolute Error (MAE) on the test set for dGNN models trained from scratch and those using transfer learning at different training sizes. For the smallest dataset size (10%), the scratch model had an MAE of 2.055, while the transfer learning model had a higher MAE of 2.502. As the dataset size increased, the scratch models continued to do better than the transfer learning models. At 40% dataset size, the scratch model's MAE was 0.499, compared to the transfer model's MAE of 0.529. For the full dataset (100%), the scratch model achieved an MAE of 0.493, while the transfer learning model had an MAE of 0.512.

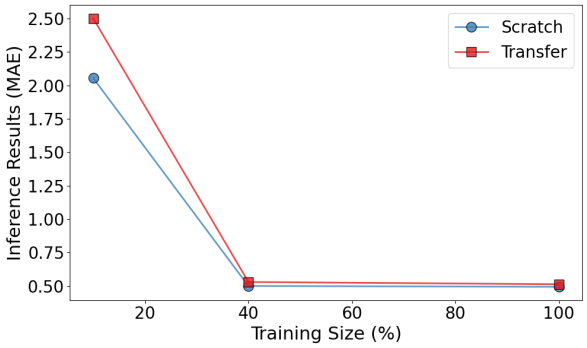

Figure 1. Comparison of Mean Absolute Error (MAE) vs. Training Size for Scratch and Transfer Learning models. The plot shows that models trained from scratch consistently performed better than those using transfer learning across all dataset sizes. The MAE for the scratch model decreases significantly as the training size increases, while the transfer learning model shows a higher MAE across all training sizes, indicating the ineffectiveness of transfer learning in this context.

The underperformance of the transfer learning models could be due to several reasons. One possibility is that the pretrained model on pristine crystal formation energies did not capture the specific details and variations present in defect structures, leading to poor fine-tuning on the defect dataset. Additionally, the transfer learning process might have introduced biases from the pretrained model that were not aligned with the characteristics of the defect data. Other common reasons for the failure of transfer learning include overfitting to the smaller defect dataset during fine-tuning, which can happen if the pretrained model is too complex and does not generalize well to the defect data. The difference between the pristine and defect datasets might have also contributed to the ineffectiveness of the transfer learning approach.

We experimented with different pooling functions (baseline dGNN pooling, dropout-styled pooling, average pooling, and max pooling) during the pretraining phase on pristine crystal formation energies. Figure 2 shows the performance comparison of these pooling functions. Among these, the average pooling function consistently delivered the best results, highlighting its effectiveness in aggregating node features for energy predictions.

The results suggest that the transfer learning approach, using a model pretrained on pristine crystal formation energies, may not be directly applicable to defect datasets without further refinement. Future work will explore the implementation of advisor-expert model

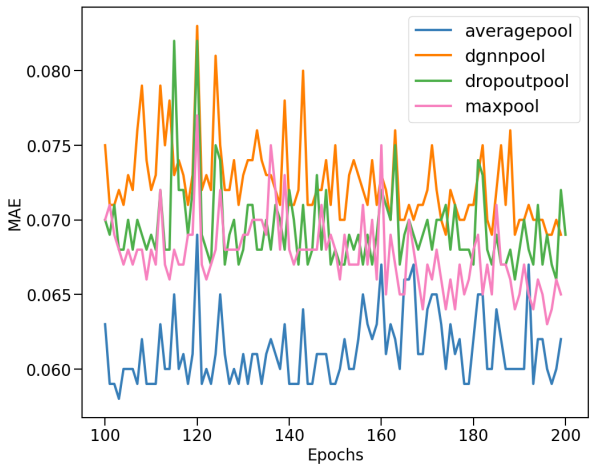

Figure 2. Performance comparison of different pooling functions (baseline dGNN pooling, dropout-styled pooling, average pooling, and max pooling) during the pretraining phase for Scratch models. The average pooling function consistently delivered the best results, showing its superior ability to aggregate node features effectively for predicting energy formations. This suggests that average pooling captures more relevant information from the graph structures compared to other pooling methods.

structures and the use of supercell representations to better capture the complexities of defect structures, aiming to refine these predictions and improve overall accuracy. Additionally, strategies to prevent overfitting and better align pretraining tasks with fine-tuning goals will be considered to enhance the effectiveness of transfer learning in this context.

To achieve better results with transfer learning in future studies, several strategies can be considered. Pretraining the model on a dataset that is more closely related to defect structures, rather than pristine crystal formation energies, can help the model learn features that are more relevant to the target task. Applying regularization techniques such as dropout or early stopping during the fine-tuning phase can prevent overfitting and improve generalization to the defect dataset. Increasing the size and diversity of the defect dataset through data augmentation techniques can help the model learn a broader range of defect characteristics and improve its predictive accuracy. Experimenting with different fine-tuning strategies, such as gradually unfreezing layers, can better adapt the pretrained model to the defect dataset without losing important pretrained features. Future research can enhance the performance of transfer learning models for defect formation energy predictions, contributing to advancements in materials science and quantum computing applications.

IV. CONCLUSIONS

In this study, we explored the effectiveness of transfer learning in enhancing the predictive capabilities of Defect Graph Neural Networks (dGNNs) for defect formation energy predictions in crystal structures. By comparing models trained from scratch with those leveraging transfer learning across different dataset sizes, we found that the scratch models consistently outperformed the transfer learning models. This underperformance suggests that the pretrained model on pristine crystal formation energies did not effectively capture the nuances of defect structures, possibly due to overfitting, biases, and differences between the pristine and defect datasets.

Despite the setbacks with transfer learning, our experiments with various pooling functions during the pretraining phase indicated that average pooling consistently provided the best results. This highlights the importance of effective feature aggregation in improving model performance.

Looking forward, our future work will focus on refining the transfer learning approach. This includes pretraining on datasets more closely related to defect structures, applying regularization techniques to prevent overfitting, and experimenting with advanced fine-tuning strategies. Additionally, we aim to explore hybrid models and incorporate adversarial training to enhance robustness and generalization capabilities.

By addressing these challenges and exploring new strategies, we hope to improve the performance of transfer learning models for defect formation energy predictions, thereby advancing materials science and contributing to the development of more efficient quantum computing applications.

V. FUTURE WORKS

Additionally, we are working to add more optimization methods to further select the best hyperparameters for model performance. In particular, we plan to expand our collection of optimizers to include genetic optimization and simulated annealing. Genetic optimization will be implemented by randomizing an initial vector of parameters, then creating new vectors changing half the values of the parameters in order to simulate a mutation, and then choosing the best performing one out of those. Simulated annealing chooses initial parameters and then continuously alters them slightly in a random direction, choosing to keep the new parameters based on a probability function involving the models' performance as well as the time elapsed. This will be favorable because unlike gradient descent, where the loss graph might get stuck at a local minimum, this

allows the model to be flexible with moving towards a global minimum, or the absolute best loss.

Additionally, we plan to add a linear layer to the model that does a linear transformation to change the sizes of parameters so that different pooling functions or pretrain/train combinations with different parameters can still be used together. Additionally, we plan to implement an Advised Experts approach, where other properties of a material, such as band gap, will be taken into account to train the defect formation energy of the material, giving us a more holistic view of factors that influence a property.

Future work will explore the implementation of advisor-expert model structures and the use of supercell representations to better capture the complexities of defect structures, aiming to refine these predictions and improve overall accuracy. Additionally, strategies to prevent overfitting and better align pretraining tasks with fine-tuning goals will be considered to enhance the effectiveness of transfer learning in this context.

In future work on this project, several approaches can be pursued to enhance the effectiveness of transfer learning. One approach is to pretrain the model on datasets that are more closely aligned with defect structures, rather than on pristine crystal formation energies, ensuring the features learned are more applicable to the target task. Another strategy is to employ regularization techniques like dropout or early stopping during fine-tuning, which can mitigate overfitting and enhance the model's ability to generalize to new data. Fine-tuning strategies could also be refined; for instance, gradually unfreezing layers of the pretrained model may help retain valuable pretrained features while adapting to new data. By pursuing these strategies, future research can significantly advance the field, contributing to both materials science and quantum computing applications.

ACKNOWLEDGEMENTS

This research was conducted at the University of North Texas (UNT) as part of a summer research program for undergraduates. We are deeply grateful to Dr. Yuanxi Wang for his invaluable guidance and support throughout this study. His expertise and insights were instrumental in shaping our research and achieving our objectives.

We would like to thank the National Science Foundation for their financial support through Award Abstract No. 2244259, REU Site: Beyond Language: Training to Create and Share Vector Embeddings across Applications. Special thanks to the REU program at UNT for providing a collaborative and enriching environment that greatly benefited this research.

[1] J. Wang, X. Zhao, G. Zou, L. Zhang, S. Han, Y. Li, D. Liu, C. Fernandez, L. Li, L. Ren, and Q. Peng. Crystal-defect engineering of electrode materials for energy storage and conversion. *Materials Today Nano*, 22:100336, 2023.

[2] R. Afre, N. Sharma, M. Sharon, and M. Sharon. Transparent conducting oxide films for various applications: A review. *REVIEWS ON ADVANCED MATERIALS SCIENCE*, 53(1):79–89, 2018.

[3] G. Zhang, Y. Cheng, J. Chou, and A. Gali. Material platforms for defect qubits and single photon emitters, 2020.

[4] A. Deml, V. Stevanovic, A. Holder, R. O'Hayre., and C. Musgrave. Intrinsic material properties dictating oxygen vacancy formation energetics in metal oxides. *Journal of Physical Chemistry Letters*, 6(10):1948–1953, 2015.

[5] N. Frey, D. Akinwande, D. Jariwala, and V. Shenoy. Machine learning-enabled design of point defects in 2d materials for quantum and neuromorphic information processing. *ACS Nano*, 14(10):13406–13417, 2020.

[6] D. Xue, P. Balachandran, J. Hogden, J. Theiler, D. Xue, and T. Lookman. Accelerated search for materials with targeted properties by adaptive design. *Nature Communications*, 7, 2016.

[7] M. Witman, A. Goyal, T. Ogitsu, A. McDaniel, and S. Lany. Defect graph neural networks for materials discovery in high-temperature clean-energy applications. *Nature Computational Science*, 3(8), 8 2023.

[8] Tian Xie and Jeffrey C. Grossman. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Phys. Rev. Lett.*, 120:145301, Apr 2018.

[9] Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, et al. Commentary: The materials project: A materials genome approach to accelerating materials innovation. *APL materials*, 1(1), 2013.

[10] Matthew Witman, Anuj Goyal, Tadashi Ogitsu, Anthony McDaniel, and Stephan Lany. A database of vacancy formation enthalpies for materials discovery, June 2023.