Response Letter

Paper Decision

Decision   Accept (Spotlight)

RESPONSE to Decision   The authors sincerely thank the Area Chair for accepting our paper and selecting it as a Spotlight talk. It is a great honor for us who were the first time submitters to any NeurIPS Workshops.

Comment   Based on the feedback by the reviewers and the recommendation by the Area Chair, the AI4Mat-NeurIPS-2024 Program Committee decided to accept this submission. Please prepare a camera-ready submission taking into account the feedback from reviewers and the meta-review.

Meta Review Comment   This submission proposes a novel approach that integrates fragment-based graph neural networks (FBGNN) with many-body expansion (MBE) theory to predict potential energy surfaces, providing a valuable contribution to computational material science. The combination of machine learning techniques with quantum mechanical methods like DFT (Density Functional Theory) and MP2 allows for higher-order corrections and offers computational savings, which are especially useful for large molecular systems.

Reviewers generally praised the work for its innovative use of FBGNN to handle many-body interactions and augment traditional first-principles methods. This hybrid approach could address scalability challenges that arise when applying DFT to large-scale systems. However, reviewers noted the need for further clarification on how fragments are chosen for first-order calculations and how the model compares to full DFT/MP2 on more complex systems. Providing explicit wall-time comparisons instead of percentage savings would also strengthen the case for the practical utility of the method.

Concerns about benchmarking against more advanced models, such as SE3Set and MACE, were raised, and additional analysis on the transferability of the method to other molecular systems was suggested. Nevertheless, the contribution is significant, and the framework holds promise for advancing AI-guided material characterization.

RESPONSE to Comment and Metal Review Comment   The authors thank the Area Chair for their suggestions and thank all Scientific Reviewers for their comments. We prepared a camera-ready version of the paper and submitted on 11/1/2024. In this version we tried our best to address all review comments and have provided our response on OpenReview. Unfortunately, due to some family emergency of the primary corresponding author and the health issue of the leading author, as well as the narrow time window for the revision, we were not able to to fully address some comments, especially the comparison of the model performance with all frontier GNN-based approaches and the holistic discussion of fragmentation strategies. However, we will continue working on these comments in the near future and update the results in our oral and poster presentations at the Workshop in anything develops.

Reviewer qoPF

Summary   The authors utilize GNN to provide higher order correction to 1-body MP2/DFT energies. Authors utilize multi-scale hierarchical graphs to effectively model many body interactions.

RESPONSE to Summary   The authors thank the Reviewer for making this summary.

Quality   The idea to use ML for higher order correction seems bit counter-intuitive at first, as one would assume that ML would be more effective in modeling largest magnitudes of energy. However it is a significant result if such correction can be cheaply used to augment first-principles energies/forces. The manuscript is well written and clearly states the goals. The authors might consider adding explanations on how do you decide on the fragments for 1st-order calculations. Also can the authors provide a specific example on where this increased accuracy of MBE correction would make a significant difference, as opposed to just using single fragment MP2.

RESPONSE to Quality   The authors thank the Reviewer for the positive feedback and the two suggestions.

We will address the Reviewer's first suggestion about the fragmentation strategies here. Reviewer aFUJ asked the same question, so we have copied our response to them as well.

At the current development stage of FBGNN-MBE, the fragmentation strategy was determined case-by-case and depended on the chemical properties the systems in question. We currently focused on the total ground state energy of a system and will investigate the excited state energy soon. We have so far implemented a "top-down" fragmentation strategy with two principles: (1) We maintain the smallest functionally meaningful unit. (2) We break only single bonds or non-bonding interactions.

In the present proof-of-concept work we studied water aggregates, phenol aggregates, and water–phenol mixtures. In these systems every single molecule was a natural fragment and an intermolecular interaction like hydrogen bond and van der Waals force plays an essential role in n-body energies ($n \geq 2$). For our future studies about organic polymers, we have tentatively planned to extend the fragmentation strategy by treating each monomer as a fragment, cleaving the carbon–carbon ($C-C$) bonds, and considering each solvent molecule as a fragment too (if any). We will try our best not to break any complete functional groups like phenyl ($-C_6H_5$) and carboxylic acid ($-COOH$) or known unit like an amino acid and a DNA base.

For a new system, we will run a low-level first-principles MBE_2 calculation where the total energy is truncated at the two-fragment interactions to make sure the fragmentation does not introduce huge estimated error to damage the chemistry. In this way the system will maintain the smallest functional or repeating unit such as a full π-conjugation, and our model can capture short-range and long-range interactions to simulate the real chemical systems, such as covalent bonds, ionic bonds, hydrogen bonds, London dispersions, dipole–dipole interactions, π–π stacking, and solvation effects.

If a system does not exhibit natural or obvious fragments, such as polyacetylene and polyethlyene, we will consider an alternative "bottom-up" strategy, where we will construct fragments from individual atoms, functional groups, or monomers until some convergence is reached using a low-level first-principles MBE_2 calculation. In this way we will allow fragments to grow around each unit and accommodate underlying effects from local electronic environment and molecular conformations.

We foresaw great room for us to explore and validate our fragmentation strategies and enhance our model versatility and accuracy. However, due to the time constraints, we were unfortunately not able to perform a systematic discussion in the present study. However, we did validated our top-down fragmentation strategy using a series of water clusters H_2O_n ($n = 7, 10, 13, 16,$ and 21) by showing first-principles $MP2\text{-}MBE_2$ and $MP2\text{-}MBE_3$ results in Tables 5-7 in Section A.5. When we treated each water molecule as a single fragment, we reached an average of 3.00% and 0.39% relative errors for these water clusters using $MP2\text{-}MBE_2$ and $MP2\text{-}MBE_3$, confirming a favorable choice of fragmentation strategy.

We have included discussions about the fragmentation strategy in Appendix A.2.

We will address the Reviewer's second suggestion about the necessity of performing MBE here. We were actually a bit confused what the Reviewer meant by "single fragment MP2".

If the Reviewer meant treating the whole aggregate as a single fragment and calculate its total energy directly using MP2, we found that the calculation was very time-consuming and became unfeasible for a larger system. For example, for random water clusters (H_2O_n) with $n = 7, 10, 13, 16,$ and 21, as given in Table 7 in Section A.5, the CPU times for running a single-point MP2 calculation using the aug-cc-pVDZ basis set were 1562.94, 18914.93, 128084.39, 51249.38, and 121046.76 seconds,

respectively. This computational complexity was estimated to follow N^4, which quickly becomes infeasible for large systems.

If the Reviewer meant treating the total energy of the whole aggregate as the sum of the MP2-calculated energy from any isolated single fragment without considering any many-fragment interactions, we found such an MP2-MBE$_1$ method introduced huge errors because these many-fragment interactions, even as weak as van der Waals forces or hydrogen bonds, played essential roles in chemical properties. For example, we examined the same series of water clusters using MP2-based MBE without any approximations from FBGNN, and found that the truncation errors with MP2-BE$_1$ were as high as 63.13, 103.19, 140.18, 180.09, and 247.90 kcal/mol. However, when we considered two-fragment interactions (MP2-MBE$_2$), the errors were reduced to 12.53, 22.98, 29.01, 38.92, and 12.53 kcal/mol. A further inclusion of three-fragment interactions (MP2-MBE$_3$ further reduced the errors to 1.36, 2.77, 3.15, 4.03, and 10.87 kcal/mol. In chemistry, we hope the errors can be as small as 1 kcal/mol, which is referred to as the chemical accuracy, so that MP2-MBE$_3$ provides a more realistic descriptions of an actual chemical system.

Questions   Can the authors comment on transferability of this method, i.e. can their model retain its accuracy on a water-cresol simulation? Instead of giving t_{save}, can the authors also provide actual wall times in doing MP2/DFT MBE calculation, vs 1 body MP2 + higher order GNN?

RESPONSE to Questions   The authors thank the Review for raising the question of transferability and suggesting about the providing wall times in doing MP2/DFT MBE calculations.

Regarding the transferability, the authors agree with the Reviewer that it is important to test our FBGNN-MBE model over diverse chemical systems with different characters of local and global graphs. Therefore in the present study we investigated the performance of FBGNN-MBE on molecular aggregates with different levels of intermolecular interactions in this study, such as water clusters with hydrogen bonds (stronger), phenol clusters with hydrogen bonds and van der Waals forces (weaker), and their mixtures. From these three systems we can preliminarily confirm about the transferability of FBGNN-MBE, as long as the data set is of high quality and diversity. However, the methyl group ($-CH_3$) did not exist in our training set, and implementation of a transfer learning strategy is beyond the scope of present study due to the time constraint. Therefore we were not very confident about the transferability to the water–cresol mixture for the moment. Having said that, we will construct more diverse data sets, design more flexible local and global graph structures, introduce semi-empirical electronic structures in the molecular descriptors [6, 7], and implement the transfer learning strategy in future studies [2], so as to enhance the transferability our FBGNN-MBE.

Regarding the comparison of wall times, we understand that the Reviewer is interested learning about the physical times we spent on a large MP2 or DFT calculation and the sum of all single-fragment MP2 or DFT calculations and MBE-estimated corrections. However, we believe that wall times were somewhat counterintuitive because our studies were performed on different computing nodes. We calculated MP2 or DFT energies for both single fragments and whole clusters on CPU nodes, using one CPU cores for each calculation, respectively. We then predicted two- and three-fragment energies based on FBGNN-MBE on GPU nodes, using four NVIDIA A100 GPU cores for each generation. Therefore we believe a more standardized comparison should occur between total CPU and GPU times. In Table 2 and Section 4.1 of the camera-ready version, we included these statistics.

References

[1] Federico Zahariev, & M. S. Gordon. "Combined quantum Monte Carlo–effective fragment molecular orbital method: fragmentation across covalent bonds." Physical Chemistry Chemical Physics 23.26 (2021): 14308-14314.

[2] David Buterez, Jon Paul Janet, Steven J. Kiddle, Dino Oglic & Pietro Lió "Transfer learning with graph neural networks for improved molecular property prediction in the multi-fidelity setting." Nature Communications 15.1 (2024): 1517.

Reviewer aFUJ

Originality   The idea of combining the fragment-based idea with the many-body expansion (MBE) theory is very good. Based on the idea of MBE, the idea of using MXMNet and PAMNet to process

2b and 3b is very good. This provides a way to reproduce high-precision PES for larger systems. The author has a thorough literature review of MBE and fragment-based GNN. Although the work is based on the fragmentation method based on previous research [1][2], this combination is worthy of recognition.

[1] Shuo Zhang, Yang Liu, & Lei Xie. "Molecular mechanics-driven graph neural network with multiplex graph for molecular structures." arXiv preprint arXiv:2011.07457, (2020).

[2] Shuo Zhang, Yang Liu, & Lei Xie. "A universal framework for accurate and efficient geometric deep learning of molecular systems." Scientific Reports, 13(1), 19171, (2023).

RESPONSE to Originality   The authors thank the Reviewer for this positive comment. However, we also hope to clarify our method a little bit to avoid confusion. We borrowed MXMNet and PAMNet from studies [1] and [2] as our backbone FBGNN approaches, so that we can say our fragmentation strategies of the whole graphs came from MXMNet and PAMNet. However, we borrowed our fragmentation strategies of chemical systems from earlier fragment-based methods like first-principles MBEs [4], such as ASMF [5] from Bettens and coworkers, and GMBE [6, 7] and energy-screened MBE [8] from Herbert and coworkers.

Quality   The implementation of this idea and the demonstration of the results are clearly presented. In the experimental analysis on the $\tilde{2}0$ and phenol systems, the discussion of the results and model performance of different models and different data sets is sufficient. In addition, the further analysis of the impact of the molecular density on the character of the electronic structures and the sensitivity analysis are worthy of recognition.

RESPONSE to Quality   The authors thank the Reviewer for this detailed analysis.

Clarity   The writing is clear and concise, and the presentation of data and program implementation is well structured. The flow of the paper is logical, and the figures and tables effectively support the arguments. Although the author discusses the weakness of the work in the analysis, it is recommended to make a brief summary in the conclusion.

RESPONSE to Clarity   The authors thank the Reviewer for the suggestion of adding a brief summary of the weakness in the Conclusion section because this revision will strengthen our concluding remarks. In the camera-ready version, we included the following sentences in Section 5: "While our FBGNN-MBE framework was promising in both efficiency and accuracy, it exhibited at least two limitations that required our attentions in future developments. First, our training sets were created using $2\times$ molecular densities, which allowed us to sample high-energy configurations and enhance ergodicity and diversity of the data set, but also introduced biases towards repulsive 2B interactions [**doi:10.1021/acs.chemrev.0c01111**]. Second, the present study focused on molecular aggregates with weak to moderate many-fragment interactions and the current fragmentation strategy did not cleave chemical bonds, so that the model transferability remained elusive. Finally, the present study did not implement the predictions for potential energy gradients (forces) or excited state energies."

Significance   While the fragmentation approach relies on previously reported methods such as MXMNet, and the experiment part only focuses on specific molecular systems H2O and phenol, this attempt is an inspiration for people to explore high-precision PES of more complex molecular systems. However, it would strengthen this paper if a wider range of fragmentation strategies were explored, potentially expanding the scope and impact of this work.[3]

[3] Unke, O. T., Stöhr, M., Ganscha, S., Unterthiner, T., Maennel, H., Kashubin, S., ... & Müller, K. R. (2024). Biomolecular dynamics with machine-learned quantum-mechanical force fields trained on diverse chemical fragments. Science Advances, 10(14), eadn4397.

RESPONSE to Significance   The authors thank the Reviewer for making the constructive feedback about the exploration of the fragmentation strategies. Reviewer qoPF raised the same concern in their comments, so we have copied our response to them here.

At the current development stage of FBGNN-MBE, the fragmentation strategy was determined case-by-case and depended on the chemical properties the systems in question. We currently focused on the total ground state energy of a system and will investigate the excited state energy soon. We have

so far implemented a "top-down" fragmentation strategy with two principles: (1) We maintain the smallest functionally meaningful unit. (2) We break only single bonds or non-bonding interactions.

In the present proof-of-concept work we studied water aggregates, phenol aggregates, and water–phenol mixtures. In these systems every single molecule was a natural fragment and an intermolecular interaction like hydrogen bond and van der Waals force plays an essential role in n-body energies ($n \geq 2$). For our future studies about organic polymers, we have tentatively planned to extend the fragmentation strategy by treating each monomer as a fragment, cleaving the carbon–carbon ($C-C$) bonds, and considering each solvent molecule as a fragment too (if any). We will try our best not to break any complete functional groups like phenyl ($-C_6H_5$) and carboxylic acid ($-COOH$) or known unit like an amino acid and a DNA base.

For a new system, we will run a low-level first-principles MBE_2 calculation where the total energy is truncated at the two-fragment interactions to make sure the fragmentation does not introduce huge estimated error to damage the chemistry. In this way the system will maintain the smallest functional or repeating unit such as a full π-conjugation, and our model can capture short-range and long-range interactions to simulate the real chemical systems, such as covalent bonds, ionic bonds, hydrogen bonds, London dispersions, dipole–dipole interactions, $\pi-\pi$ stacking, and solvation effects.

If a system does not exhibit natural or obvious fragments, such as polyacetylene and polyethlyene, we will consider an alternative "bottom-up" strategy, where we will construct fragments from individual atoms, functional groups, or monomers until some convergence is reached using a low-level first-principles MBE_2 calculation. In this way we will allow fragments to grow around each unit and accommodate underlying effects from local electronic environment and molecular conformations.

We foresaw great room for us to explore and validate our fragmentation strategies and enhance our model versatility and accuracy. However, due to the time constraints, we were unfortunately not able to perform a systematic discussion in the present study. However, we did validated our top-down fragmentation strategy using a series of water clusters H_2O_n ($n = 7, 10, 13, 16$, and 21) by showing first-principles $MP2\text{-}MBE_2$ and $MP2\text{-}MBE_3$ results in Tables 5-7 in Section A.5. When we treated each water molecule as a single fragment, we reached an average of 3.00% and 0.39% relative errors for these water clusters using $MP2\text{-}MBE_2$ and $MP2\text{-}MBE_3$, confirming a favorable choice of fragmentation strategy.

We have included discussions about the fragmentation strategy in Appendix A.2.

References

[4] Michael A. Collins & Ryan P. A. Bettens. "Energy-based molecular fragmentation methods." Chemical Reviews 115.12 , (2015): 5607–5642.

[5] Isabel Rodrigo, Sofía Caruncho, Concepción Alonso, Antonio Gómez-Outes & Barbara Mulloy. "Low Molecular Weight Heparins, Biological Drugs close to Non-Biological Complex Drugs." Non-Biological Complex Drugs: The Science and the Regulatory Landscape (2015): 291–331.

[6] Ryan M. Richard & John M. Herbert. "A generalized many-body expansion and a unified view of fragment-based methods in electronic structure theory." The Journal of Chemical Physics 137.6 (2012): 064113.

[7] Ryan M. Richard & John M. Herbert. "Many-body expansion with overlapping fragments: Analysis of two approaches." Journal of Chemical Theory and Computation 9.3 (2013): 1408–1416.

[8] Kuan-Yu Liu & John M. Herbert. "Energy-screened many-body expansion: A practical yet accurate fragmentation method for quantum chemistry." Journal of Chemical Theory and Computation 16.1 (2019): 475–487.

[9] Oliver T. Unke, Stefan Chmiela, Huziel E. Sauceda, Michael Gastegger, Igor Poltavsky, Kristof T. Schütt, Alexandre Tkatchenko, & Klaus-Robert Müller "Machine learning force fields." Chemical Reviews 121.16 (2021): 10142-10186.

[10] Federico Zahariev & M. S. Gordon. "Combined quantum Monte Carlo–effective fragment molecular orbital method: fragmentation across covalent bonds." Physical Chemistry Chemical Physics 23.26 (2021): 14308-14314.

Reviewer PGtA

Summary   The study aims to predict higher-order energy terms (specifically 2-body and 3-body interactions) based on many-body expansion theory to achieve a more accurate quantitative understanding of molecular properties. The authors generated a new dataset by conducting MD simulations for the 2-body and 3-body energy tasks. They utilized MXMNet and PAMNet as backbone models and compared their performance with other existing models.

RESPONSE to Summary   The authors thank the Reviewer for making an exact summary of our study.

Strengths

1. The work contributes a new dataset, and the model demonstrates good fitting performance on the data.

2. The authors provide a comprehensive analysis of the results from various perspectives.

RESPONSE to Strengths   The authors thank the Reviewer for summarizing the strengths of our study.

Weaknesses

1. I have concerns about the comparison with other models used to claim state-of-the-art performance, as these models are either outdated or lack higher-order terms for this task. More advanced models, such as SE3Set and MACE with fragment-based representations or many-body terms, would be more appropriate for comparison.

2. The model's novelty appears limited. Although MXMNet-MBE and PAMNet-MBE are employed for the tasks, there are no clear descriptions of how they differ from the original MXMNet and PAMNet models. The modifications made to the compared models are not described in sufficient detail, which may hinder replicability in the current manuscript.

RESPONSE to Weaknesses   The authors thank the Reviewer for raising the two weaknesses of our study and providing corresponding constructive feedbacks.

We will address the Reviewer's first suggestion about adding comparisons to more advanced models here. We agree with the Reviewer that a comparison with more up-to-date models is more appropriate for the present study. Therefore we are actively working on the implementation of SE3Set [1], MACE [2], MGSSL [3], MolGAT [4], FragGen [5], FragNet [6], GraphFP [7], Subgraphormer [8], and ESC-GNN [9] in our study. The authors of SE3Set have not published their source codes because their paper is still under review with another conference (11/1/2024), and we are still waiting for their response. Due to the time constraint, the family emergency of the corresponding author and the health issue of the leading author, we are very sorry to say that we were not able to implement the rest of these models by 11/1/2024. If anything develops before the Workshop, we will make sure to include them in our oral and poster presentations.

We will address the Reviewer's second suggestion here by providing more information about the fragment-based implementation of our models and the pseudo-algorithms to facilitate the reproduction of our results.

Regarding the novelty of the fragment-based approach, we believe that our major novelty did not lie in the FBGNN models themselves because we simply borrowed MXMNet and PAMNet from the work of Xie and coworkers as our backbone GNN models. Instead, our novelty lied in the adoption of FBGNN models to the fragment-based QM approaches such as MBE theory. In this way, we aligned the subgraph-based hierachies of GNNs with the chemical hierarchies of complex molecular systems, broke down such a system into M chemically meaningful pieces, and utilized a "divide-and-conquer" strategy rather than calculating the entire system in one shot. This treatment was estimated to reduce the computational cost from N^3M^3 to NM^3 (DFT) or from N^5M^5 to N^3M^5 (MP2) for a system containing N fragments with M basis functions each fragment.. The Reviewer can refer to our responses to Reviewers qoPF and aFUJ for a discussion of our fragmentation strategies, which were

different from the original works of MXMNet and PAMNet and provided a physical understanding of how molecular fragments work together towards the chemical properties.

Regarding the replicability of our FBGNN-MBE models, we have provided the pseudo-algorithms that outline the protocols in MXMNet-MBE and PAMNet-MBE in Section A.1 of the camera-ready submission.

Summary   Given these reasons, I believe this paper falls below the acceptance threshold.

RESPONSE to Summary   The authors thank the Reviewer for making this suggestion even it disfavors our acceptance. We acknowledge all of our weakenesses and have been trying our best to improve the present study.

References

[1] Hongfei Wu, Lijun Wu, Guoqing Liu, Zhirong Liu, Bin Shao, & Zun Wang. "SE3Set: Harnessing equivariant hypergraph neural networks for molecular representation learning." arXiv preprint arXiv:2405.16511 (2024).

[2] Ilyes Batatia, Dávid Péter Kovács, Gregor N. C. Simm, Christoph Ortner, & Gábor Csányi. "MACE: Higher order equivariant message passing neural networks for fast and accurate force fields." Advances in Neural Information Processing Systems 35 (2022): 11423–11436.

[3] Zaixi Zhang, Qi Liu, Hao Wang, Chengqiang Lu,& Chee-Kong Lee. "Motif-based graph self-supervised learning for molecular property prediction." Advances in Neural Information Processing Systems 34 (2021): 15870–15882.

[4] Mesfin Diro Chaka, Chernet Amente Geffe, Alex Rodriguez, Nicola Seriani, Qin Wu, & Yedil-fana Setarge Mekonnen. "High-Throughput Screening of Promising Redox-Active Molecules with MolGAT." ACS omega 8.27 (2023): 24268–24278.

[5] Mengyao Yu, Yundian Zeng, Mingyang Wang, Chenqing Hua, Sunliang Cui, Peichen Pan, Chang-Yu Hsieh, Tingjun Hou, Odin Zhang, Yufei Huang, Shicheng Chen, Xujun Zhang, Haitao Lin, Zhenxing Wu, Huifeng Zhao, Zaixi Zhang & Yu Kang. "FragGen: Towards 3D Geometry Reliable Fragment-based Molecular Generation." In press in Chemical Science (2024).

[6] Gihan Panapitiya, Peiyuan Gao, C Mark Maupin, & Emily G Saldanha. "FragNet: A Graph Neural Network for Molecular Property Prediction with Four Layers of Interpretability." arXiv preprint arXiv:2410.12156 (2024).

[7] Kha-Dinh Luong and Ambuj Singh. "Fragment-based pretraining and finetuning on molecular graphs." Advances in Neural Information Processing Systems 36 (2024): 17584–17601.

[8] Bar-Shalom, Guy, Beatrice Bevilacqua, & Haggai Maron. "Subgraphormer: Unifying Subgraph GNNs and Graph Transformers via Graph Products." arXiv preprint arXiv:2402.08450 (2024).

[9] Zuoyu Yan, Junru Zhou, Liangcai Gao, Zhi Tang, & Muhan Zhang. "An Efficient Subgraph GNN with Provable Substructure Counting Power." Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (2024).

Reviewer rvQW

Comment   Authors propose FBGNN-MBE, which aim to combine the strengths of many body expansions and graph neural nets. They divide the structure into small fragments, get 1st body contributions from first-principles and 2nd and 3rd body contributions from MXMNet or PAMNet models. I find the idea to be interesting and timely, given the increasing importance of modeling largest molecular systems with higher accuracy in QM. Authors test their approach in several benchmark systems (water, phenol and water:phenol). They run MD with their new methods, and show overall that the method results in significant computational speed up at the expense of reasonable-looking errors in energies. Manuscript is well-written, clear and timely. Minor suggestion: could authors comment on the nature of forces (e.g. are they energy conserving etc.) obtained with this setup, as they would be used in MD?

RESPONSE The authors thank the Reviewer for the positive comment and the suggestions about discussing the nature of forces obtained from our model set up and their suitability for MD simulations. In the present set up, we approximated the first-principles potential energy surfaces (PESs), or in other words, geometry-dependent full-dimensional electronic energies under the Born–Oppenheimer approximation at $T = 0$ K, using FBGNN-MBE generated counterparts. We can evaluate all relevant forces as the gradients of these potential energies [1, 2]. In this way, we expected to replace a first-principles PES with a FBGNN-MBE-generated PES in *ab initio* molecular dynamics (AIMD). We are not high-level experts for MD simulations, so we are not 100% sure what the Reviewer meant by the nature of forces. However, we will try our best to answer this question. Please forgive us if any of the terminology below is inexact.

If the Reviewer asked about the underlying rules like energy-conservation for the forces, we believe that they really depended on the overall set up of the AIMD simulations rather than that of our FBGNN-MBE-generated PES because our PES can be used in an identical way to the first-principles PES in AIMD. If the AIMD simulation is set up at constant NVE (microcanonical ensemble), the energy is definitely conserved. However, if AIMD simulation is set up at constant NVT (canonical ensemble) or μVT (grand canonical ensemble), the temperature remains a constant through a computational thermostat but the energy is no longer conserved.

If the Reviewer asked about the types of the forces, we believe that they really depended on the chemical and physical properties of the system in question. In the present proof-of-concept study, the systems were simple molecular aggregates, so that the forces within each fragment came from chemical bonds, and the forces between fragments came from hydrogen bonds and van der Waals interactions. In our future studies, the forces between fragments may also come from ionic bonds, covalent bonds, and metallic bonds. We will tailor our FBGNN-MBE models and fragmentation strategies based on the types of the forces. To improve the capacity of physical interpretation of FBGNN-MBE, we also plan to implement the idea of energy decomposition analysis (EDA) so as to fraction the interaction energies further into different types of attractive and repulsive terms, such as electrostatics, Pauli repulsion, exchange–correlation, and dispersion [3]. By implementing EDA to our PES and training for all components, we expect to inject physical information to FBGNN-MBE and enhance the model understanding of the physical behaviors of complex organic systems. In the camera-ready submission, we have included our discussion about the nature of the forces in the Sections 5 and A.5

[1] Joseph P. Heindel, Kristina M. Herman, & Sotiris S. Xantheas. "Many-body effects in aqueous systems: Synergies between interaction analysis techniques and force field development." Annual Review of Physical Chemistry 74.1 (2023): 337-360.

[2] Oliver T. Unke, Stefan Chmiela, Huziel E. Sauceda, Michael Gastegger, Igor Poltavsky, Kristof T. Schütt, Alexandre Tkatchenko, & Klaus-Robert Müller "Machine learning force fields." Chemical Reviews 121.16 (2021): 10142-10186.

[3] Lili Zhao, Moritz von Hopffgarten, Diego M. Andrada, & Gernot Frenking. "Energy decomposition analysis." WIREs Computational Molecular Science 8.3 (2018), e1345.