
Accelerated Isotopologue Reduced Partition Function Ratio Prediction with Orbital-based Deep Learning

Simon Andren*

Caltech

sandren@caltech.edu

Beom Seok Kang*

Caltech

bkang@caltech.edu

William A. Goddard III

Caltech

wag@caltech.edu

John M. Eiler

Caltech

eiler@caltech.edu

Anima Anandkumar

Caltech

anima@caltech.edu

Abstract

Predictions of the reduced partition function ratios (RPFRs) of isotopologues, versions of molecules differing in position and numbers of isotopes, form a predictive framework for interpreting isotopic data from natural samples, offering insights into formation pathways and environmental conditions. However, traditional computational approaches are either computationally expensive or insufficiently accurate. Here, we employ OrbNet-Equi, a state-of-the-art orbital-based deep learning framework, speeding up predictions of RPFRs by a factor of ~ 1000 , while maintaining accuracy comparable to density functional theory (DFT). To optimize isotopic predictions, we incorporate element-wise pooling and masking strategies. OrbNet-Equi achieves target accuracy (sub-percent for ^2H and sub-permille for ^{13}C , ^{15}N , ^{18}O) with training sets as small as 500 molecules. Using the full dataset of 100,000 molecules at 300 K, the model yields a mean absolute permille error about four times smaller than the target threshold and predicts $>95\%$ of RPFRs within the desired accuracy. Compared to other non-DFT approaches, OrbNet-Equi reduces mean absolute permille error by up to 10-fold. This establishes a computational framework capable of extending RPFR predictions to reaction networks relevant to geochemical and biochemical systems.

1 Introduction

Isotopes have subtle differences in their chemical and physical properties, resulting in measurable differences in their distributions throughout natural materials. The enrichment or depletion of rare isotopes between phases and chemical species is quantified by the fractionation factor α and when coupled with isotopic measurements provides a powerful tool for reconstructing paleoclimate, tracing metabolic and environmental pathways, and interpreting forensics signatures [1–4]. Predictions of fractionation factors are important for translating isotopic observations into answers about underlying processes. Although analytical techniques have advanced substantially, expanding the range of isotopic measurements and measurable molecule [5], comparable improvements in predictive modeling of the fractionation factors are lacking, leaving a gap between measurement and interpretation [6].

The reduced partition function ratio (RPFR) provides a statistical–mechanical basis for predicting isotopic fractionation factors. Because heavier isotopes lower molecular vibrational energies, they preferentially partition into isotopologues, molecules that differ only in the number and placement of rare isotopes, that minimize the system’s free energy at equilibrium. This behavior is commonly

*Equal contribution.

quantified within the Urey–Bigeleisen–Mayer (UBM) framework, which expresses RPFRs as functions of temperature and the harmonic vibrational frequencies of the unsubstituted and isotopically substituted isotopologues. Equilibrium fractionation factors can be predicted from the RPFRs of reactions and products, and kinetic fractionation factors can be derived from those of reactants and transition states. Coupled with reaction networks, these predictions yield condition-specific (e.g., temperature-dependent) isotopic fingerprints that can be used to trace chemical and physical processes. In the UBM framework, the RPFR can be approximated as:

$$\beta = \prod_{j=1}^{\alpha} \left(\frac{\omega'_j}{\omega_j} \cdot \frac{e^{-u'_j/2}}{e^{-u_j/2}} \cdot \frac{1 - e^{-u_j}}{1 - e^{-u'_j}} \right) \quad (1)$$

where the prime symbol ($'$) denotes the isotopically substituted molecules, β is RPFR, α is the number of vibrational modes, ω_j is the j -th harmonic vibrational frequency, and $u_j = \omega_j \hbar / (k_B T)$.

Several approaches exist for obtaining RPFRs, each with distinct advantages and limitations. Density functional theory (DFT) is commonly employed to calculate the isotopologue-specific harmonic frequencies required by the UBM framework to predict equilibrium isotope effects (EIEs), with applications ranging from distinguishing biotic and abiotic methane sources to reconstructing paleotemperatures [2, 7]. However, the cubic scaling of geometry optimizations and force constant calculations with system size makes DFT prohibitively expensive for larger molecules, limiting its practical use. Alternative strategies that reduce computational cost, such as the cut-off method [8] or the fully empirical Galimov bond additivity model [9], compromise accuracy and therefore fall short for reliable RPFR predictions.

Machine learning (ML) has emerged as a powerful tool in quantum chemistry, offering a practical balance between computational cost and accuracy. A growing body of work has shown that ML models can reliably predict molecular properties with significantly reduced computational demands [10–13]. Among these, orbital learning has demonstrated notable improvements in both accuracy and data-efficiency by leveraging orbital-based features that more closely align with the underlying quantum mechanics [14–22]. Given the high cost of traditional quantum calculations, the enhanced data-efficiency of orbital learning makes it especially useful for data-scarce applications.

In this study, we utilize the state-of-the-art orbital-based deep learning method, OrbNet-Equi [15], to predict the RPFRs of singly-substituted isotopologues. Leveraging the dataset of RPFRs calculated using DFT/B3LYP [23, 24], we train models that can predict the RPFRs of deuterium (^2H)-substituted isotopologues (the *H* model) and models that can predict the RPFRs of ^{13}C -, ^{15}N -, and ^{18}O -substituted isotopologues (the *non-H* model). For this, we reformulate the singly-substituted isotopologues RPFRs prediction task to a node-wise property prediction task, and apply several task-specific operations, such as element-wise masking and element-wise decoding.

The results reveal the robust performance of OrbNet-Equi as an RPFR predictor. Only with 500 molecules, both the H and the non-H models for 300 K achieve the initially targeted accuracies, which are 10 permille (\textperthousand) for ^2H and 1% for ^{13}C , ^{15}N , and ^{18}O . Furthermore, OrbNet-Equi achieves high accuracy in RPFR prediction, with 95.1% of deuterium, 96.9% of ^{13}C , 97.4% of ^{15}N , and 95.6% of ^{18}O isotopologue predictions falling within their corresponding targeted accuracies. Additionally, the proposed framework achieves \sim 1000-times speedup compared to DFT. The notable data-efficiency and accuracy of OrbNet-Equi for the RPFR predictions highlight its robustness and utility for accelerating isotope-related applications.

2 Methods

2.1 The QM9s-RPFR Dataset

We use the QM9s-RPFR dataset, a dataset based on the QM9s dataset [23], to train OrbNet-Equi for predicting RPFRs calculated at the B3LYP/def-TZVP level of theory. The dataset will be publicly available upon the upcoming publication, with several analyses including data distribution and comparison to experimental values. The molecules represented in QM9 are all neutral closed-shell molecules, each with up to 9 heavy atoms (C, N, O, F). The RPFRs are derived from the vibrational frequencies for each isotopologue, converted into RPFRs using (1).

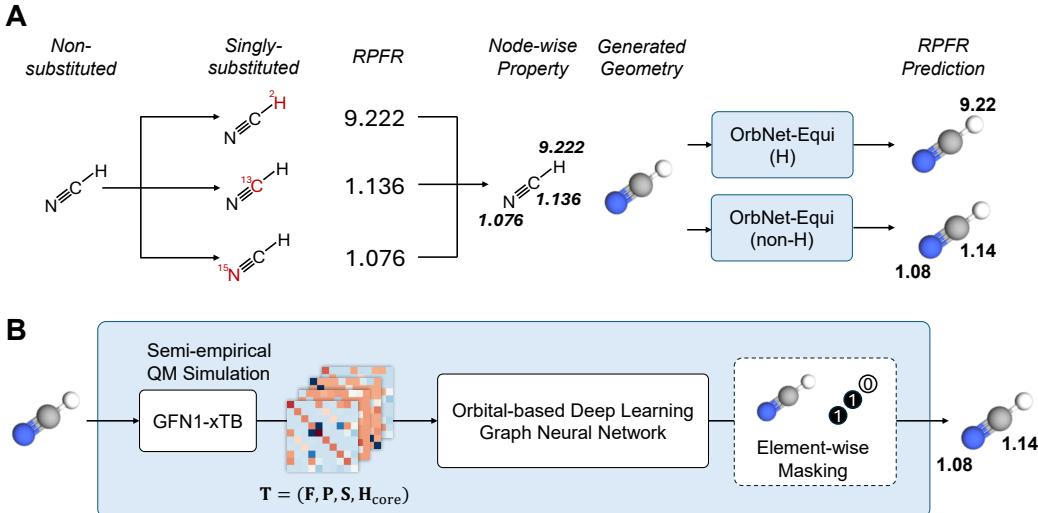


Figure 1: Overview of the project. (A) The prediction of singly substituted isotopologue RPFRs is reformulated as a node-wise property prediction task. Two OrbNet-Equi models—an H-model and a non-H model—are trained accordingly. (B) Schematic of the overall framework. Element-wise masking enables selective RPFR prediction for specific isotopic substitutions.

The dataset comprises 128,817 data points in total. We partitioned this into a training set of 100,000 molecules, a validation set of 5,000 molecules, and a test set containing the remaining molecules. To evaluate the performance of OrbNet-Equi across different training sizes, the training and validation sets were further subdivided into smaller subsets. In all cases, the size of the validation subset was fixed at 5% of the corresponding training subset. For example, a training case using 10,000 molecules was validated using 500 molecules. The number of atoms per molecule ranges from 3 to 29, meaning a single molecule can contribute up to 29 RPFR data points, which correspond to different singly-substituted isotopologues.

2.2 The OrbNet-Equi Framework

Achieving DFT-level accuracy without incurring prohibitive computational cost is essential for predictive modeling of equilibrium isotope effects in reaction networks. Traditional approaches, such as the cut-off method and the Galimov model, fall short of this requirement. To address this, we employ OrbNet-Equi, a state-of-the-art orbital-based deep learning framework that has demonstrated the highest accuracy across a wide range of quantum chemical property predictions.

The OrbNet-Equi framework [15] is a QM-informed geometric deep learning framework. The method consists of two main operations: (1) the semi-empirical quantum mechanical (SEQM) calculations using GFN1-xTB [25] to create quantum mechanical matrices (QMMs), and (2) predicting chemical properties using QMMs with an E(3) (or SE(3)) equivariant GNN.

Geometry Optimization The goal of this work is to predict RPFRs of singly-substituted isotopologues without the excessive cost of quantum chemical calculations using DFT. We intend to use OrbNet-Equi [15] to predict RPFRs, which requires the atomic coordinates information of molecules. However, in typical practical situations, only the structural information (i.e., atoms and their connectivity) is provided, and geometrical information (i.e., 3D coordinates of atoms) is missing. Therefore, a reliable and consistent method is required to generate the geometries of molecules inexpensively.

For this purpose, we employ GFN1-xTB [25], an SEQM method, to optimize the geometries of molecules inexpensively. GFN1-xTB offers an excellent balance between speed and accuracy. Several studies revealed the robust performance of GFN1-xTB in predicting chemical interactions [26, 27]

and geometry optimization [28–30]. The GFN1-xTB-level optimized geometries are then used directly as inputs to the OrbNet-Equi models.

Orbital Features GFN1-xTB performs self-consistent field (SCF) procedures to create QMMs of molecules, \mathbf{O} , which are given by:

$$(\mathbf{O})_{AB}^{\mu\nu} = \langle \phi_A^\mu | \hat{\mathcal{O}} | \phi_B^\nu \rangle, \quad (2)$$

where $\mu = (n, l, m)$ and $\nu = (n', l', m')$ are orbital indices, A and B are atom indices, n , l , and m are principal, angular, and magnetic quantum numbers, and $\hat{\mathcal{O}}$ is an operator.

The QMMs, $\mathbf{T} = (\mathbf{F}, \mathbf{P}, \mathbf{S}, \mathbf{H}_{\text{core}})$, each represents different aspects of the converged electronic structure. The Fock matrix \mathbf{F} reflects electronic energies, while the core Hamiltonian \mathbf{H}_{core} includes kinetic energy and nuclear attraction. The density matrix \mathbf{P} encodes the electronic charge distribution, and the overlap matrix \mathbf{S} describes the extent of orbital overlap. Together, these matrices characterize the electronic structures of each molecule at the GFN1-xTB level of theory. Relevant studies exhibited superior accuracy, data-efficiency, and generalization using this set of features [14–16].

The QMMs are $\text{SE}(3)$ -equivariant, where each block corresponding to the interaction between two shells with angular quantum numbers l and l' transforms predictably for a roto-translational action \mathcal{R} ,

$$(\mathcal{R} \cdot \mathbf{O})_{AB}^{l,l'} = \mathcal{D}^l(\mathcal{R})(\mathbf{O})_{AB}^{l,l'} \left(\mathcal{D}^{l'}(\mathcal{R}) \right)^\dagger, \quad (3)$$

where $\mathcal{D}^l(\mathcal{R})$ is the Wigner-D matrix of degree l , and the dagger denotes the Hermitian conjugate. This construction guarantees that the features transform consistently with the underlying group action, thereby ensuring that the overall framework preserves both rotational and translational symmetries.

To effectively predict RPFRs, we introduce modifications to the default OrbNet-Equi model. The two primary changes are: (1) element-wise decoding (pooling), and (2) element-wise masking, both of which are described below.

Element-wise Decoding To enhance OrbNet’s ability to predict RPFRs across different elements, we introduce an additional element-specific decoding layer, where the biases are initialized with the average RPFRs of each element. Because different elements exhibit distinct mass ratios relative to their heavier isotopes, this element-specific decoding helps the model to account for such variations. Together, these operations constitute the pooling mechanism. The predicted RPFRs are given by:

$$\begin{aligned} (\mathbf{h}_A^o)_{lp} &= \mathbf{W}_{lp,z_A}^o \cdot (\mathbf{h}_A^f)_{lp}, \\ \hat{y}_A &= \mathbf{W}^f \cdot \|\mathbf{h}_A^o\| + b_{z_A}, \end{aligned} \quad (4)$$

where \mathbf{h}_A^f denotes the hidden representation of atom A after the final layer, p is the parity index, and \mathbf{W}_{lp,z_A}^o is a learnable weight matrix specific to the spherical degree l , parity p , and atomic number z_A . \mathbf{W}^f is a shared learnable final pooling matrix, and b_{z_A} is an element-wise bias initialized to the mean RPFR values for each element number z_A . To reiterate, the prediction \hat{y}_A corresponds to the RPFR of the singly-substituted isotopologue in which atom A is replaced by a heavier isotope.

2.3 Target and Loss

Problem Setting We aim to predict the RPFRs of isotopologues containing ^{13}C , ^{15}N , ^{18}O , or ^2H , which are $\text{E}(3)$ -invariant scalar molecular properties. For singly-substituted isotopologues, we reformulate the RPFRs prediction as an atom-wise (node-wise) prediction task. Since the deviation in RPFR arises from substituting a specific atom with a heavier isotope, the RPFR becomes a local atomic property—one per atom per molecule. This formulation is valid under the assumption that only one heavy isotope per element is considered, which holds for our study. It significantly reduces computational cost and allows batch prediction of multiple RPFRs per molecule in a single model inference. Moreover, for elements with multiple rare isotopes (e.g., ^{17}O and ^{18}O), RPFRs associated with other isotopes can be inferred from the predicted value using mass-based scaling laws.

The experimental precision of high-dimensional isotopic measurements from natural samples is typically 1–5 % for non-hydrogen elements and >10 % for deuterium [3, 31]. Accordingly, we set target accuracies for RPFR prediction at 1 % for isotopologues singly substituted with ^{13}C , ^{15}N , or ^{18}O , and 10 % for those singly substituted with deuterium (^2H). Model performance is evaluated by reporting both the mean permille error and the fraction of species within these thresholds.

Loss Deuterium-substituted isotopologues and those containing heavier isotopes of carbon, nitrogen, or oxygen exhibit RPFRs of different magnitudes, resulting in distinct target accuracies and loss scales. To account for this, we train two separate models for RPFR prediction: a *non-H model* and an *H model*. The non-H model is designed to predict RPFRs for singly substituted isotopologues with ^{13}C , ^{15}N , or ^{18}O , while the H model targets ^2H (deuterium) substitutions. For each model, the loss function \mathcal{L} is computed selectively by applying an atom-type-specific mask: hydrogen atoms are excluded in the non-H model, and non-hydrogen atoms are excluded in the H model. This masking ensures that each model learns only from the atomic substitutions it is intended to predict. The loss functions are defined as follows:

$$\mathcal{L}_i^{\text{non-H}} = \frac{1}{\sum_A^{N_i} m_A} \sum_A^{N_i} m_A \cdot \mathcal{L}(\hat{y}_A, y_A), \quad m_A = \begin{cases} 1, & \text{if atom } A \text{ is C, N, or O,} \\ 0, & \text{otherwise,} \end{cases} \quad (5)$$

$$\mathcal{L}_i^{\text{H}} = \frac{1}{\sum_A^{N_i} m_A} \sum_A^{N_i} m_A \cdot \mathcal{L}(\hat{y}_A, y_A), \quad m_A = \begin{cases} 1, & \text{if atom } A \text{ is H,} \\ 0, & \text{otherwise,} \end{cases} \quad (6)$$

where $\mathcal{L}_i^{\text{non-H}}$ and \mathcal{L}_i^{H} are the losses of the i -th training batch in non-H and H models, N_i is the number of atoms in the batch, \hat{y}_A and y_A are the predicted and true RPFR of the isotopologue formed by substituting the atom A , and m_A is the mask. Each training batch has a fixed number of molecules but a varying number of atoms, thus the losses are formulated by averaging each atom's loss for a consistent size of gradients. Here, we define the loss function \mathcal{L} as the SmoothL1 function [32].

$$\mathcal{L}(x, y) = \text{SmoothL1}(x, y) = \begin{cases} \frac{1}{2\tau}(x - y)^2, & \text{if } |x - y| < \tau \\ |x - y| - \frac{1}{2}\tau, & \text{otherwise,} \end{cases} \quad (7)$$

and τ is a hyperparameter, the threshold for the linear-quadratic transition. The models learn the mapping between orbital features and RPFRs by minimizing the corresponding loss function. We use $\tau = 0.0001$ for training H models and $\tau = 0.01$ for training non-H models.

Metric To evaluate the performance of the models during training, validations were performed after each epoch of training. Mean absolute permille (%) error (MAPE) was used as the metric for evaluations, which is defined by:

$$\text{MAPE} = 1000 \times \frac{1}{N} \sum_A^N \left| \frac{\hat{y}_A - y_A}{y_A} \right| (\%), \quad (8)$$

where N is the number of atoms for the RPFR predictions within the evaluation set, A is the atomic index, and \hat{y}_A and y_A are the predicted RPFR and the ground-truth RPFR, respectively, of the isotopologue with a substitution on atom A . The best model with the lowest validation MAPE was taken to be evaluated with the test set.

3 Results

OrbNet-Equi achieves MAPEs of 0.23% for non-Hydrogen isotopologues and 2.75% for hydrogen isotopologues, which is about a factor of five improvement for non-H and about a factor of ten improvement for H, as shown in Table 1. Furthermore, 95% of all predictions are within the target threshold of 10%, for H substituted isotopologues and 1%, for non-H substituted isotopologues, showing improved performance compared to traditional models.

Table 1: Performance of Cut-Off (N=2), Galimov, and OrbNet-Equi models trained using 100,000 molecules for RPFRs at 300 K. Here, N denotes the cut-off (N-hop) from the atom of interest. All metrics reported are such that lower values indicate better performance.

Model	Type	MAPE (%)	Std (%)	Q95 (%)
Galimov	H	26.62	24.04	93.71
Cut-Off (N=2)	H	26.76	26.23	95.76
OrbNet-Equi (Ours)	H	2.75	7.02	9.77
Galimov	non-H	2.71	3.07	10.05
Cut-Off (N=2)	non-H	1.10	1.34	4.41
OrbNet-Equi (Ours)	non-H	0.23	0.44	0.77

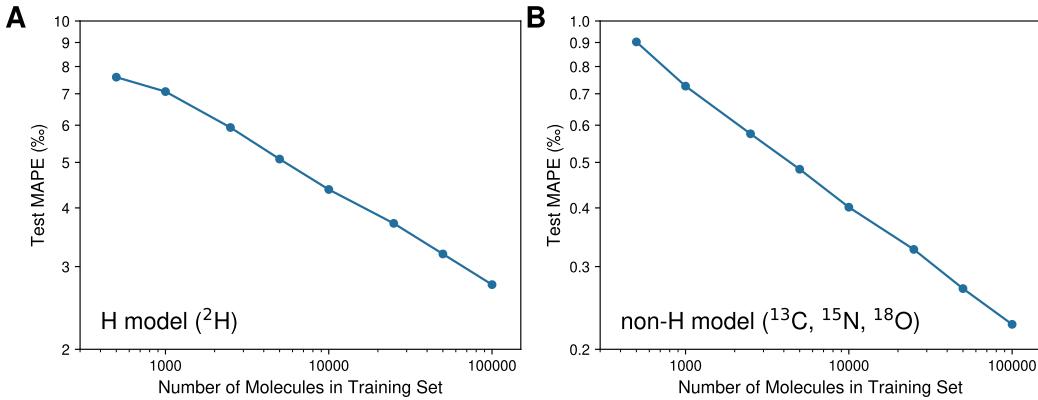


Figure 2: Learning curves of OrbNet-Equi to the 300 K RPFRs prediction. (A) Learning curve of the H OrbNet-Equi model. With only 500 molecules, it reaches the experimental accuracy (10%). (B) Learning curve of the non-H OrbNet-Equi model. With only 500 molecules, it reaches the experimental accuracy (1%).

Learning curves for the H (Figure 2A) and non-H (Figure 2B) models at 300 K show that OrbNet-Equi achieves sub-permille accuracy for non-H RPFR predictions and sub-percent accuracy for H RPFR predictions with as few as 500 training molecules. Furthermore, as the training set size increases, prediction accuracy improves further.

For both the H and the non-H models trained using 100,000 molecules at 300 K, the test error distributions exhibit bell-shaped curves centered around zero, as shown in Figure 3. Most predictions fall within error thresholds: below 10% for deuterium and below 1% for ¹³C, ¹⁵N, and ¹⁸O. Specifically, 95.1% of deuterium isotopologue predictions fall within 10%, while 96.9% of ¹³C, 97.4% of ¹⁵N, and 95.6% of ¹⁸O isotopologue predictions fall within 1%.

Additionally, we observe a higher MAPE and a lower ratio of species under the target accuracy for ¹⁸O-substituted isotopologue RPFR predictions, with the error distribution more spread than for other species, as shown in Figure 3. This higher error likely originates from the greater mass ratio relative to its lighter isotope counterpart, ¹⁶O. The greater mass ratio induces larger RPFR values, causing both the magnitude of labels and the prediction uncertainty to increase.

We investigate the temperature dependence of RPFR training, with test MAPEs of OrbNet-Equi models trained at various temperatures shown in Figure 4. Interestingly, the test MAPE exhibits a clear power-law decay with temperature, showing log-linear behavior. We explain this behavior from separate model biases in the temperature-dependent and independent terms, in Appendix A.2.

Furthermore, we compare the computation times for evaluating RPFRs using DFT and OrbNet-Equi. The analysis shows that the proposed framework achieves a mean speedup of approximately ~ 1087

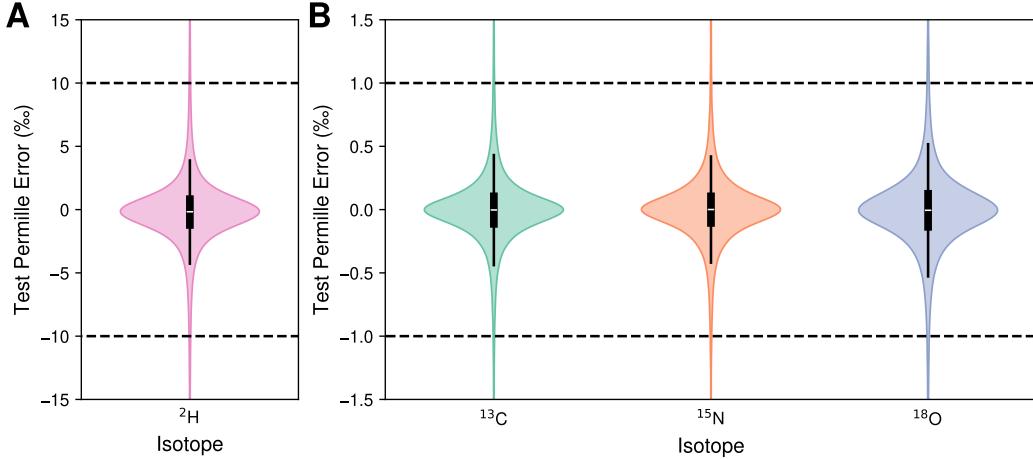


Figure 3: Test set error distribution violin plots for the RPFRs prediction at 300 K using models trained with 100,000 molecules. (A) Violin plot of the permille error distribution for ^2H isotopologues RPFRs prediction. (B) Violin plot of the permille error distribution for ^{13}C , ^{15}N , and ^{18}O isotopologues RPFRs prediction.

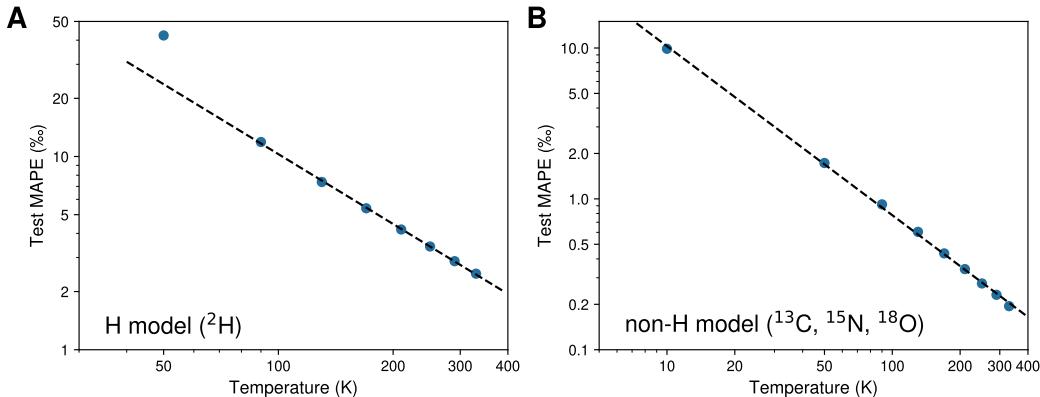


Figure 4: Temperature dependence of test MAPEs for the RPFRs prediction. All models are trained using 100,000 molecules. (A) Test set MAPEs for ^2H isotopologues RPFRs prediction at different temperatures. (B) Test set MAPEs for ^{13}C , ^{15}N , and ^{18}O isotopologues RPFRs prediction at different temperatures.

and a median speedup of approximately ~ 970 compared to DFT/B3LYP/def-TZVP. A more detailed distribution of computation times is provided in Appendix D.

4 Discussion

Performance This study demonstrates the effectiveness of OrbNet-Equi in predicting RPFRs of singly substituted isotopologues. The model is highly data-efficient, where with only 500 training molecules, it achieves mean absolute percentage errors (MAPEs) below 10% for ^2H and below 1% for ^{13}C , ^{15}N , and ^{18}O , with further reductions as the training set size increases. Remarkably, when trained using 100,000 molecules, 95.1% of deuterium isotopologue predictions fall within 10%, while 96.9% of ^{13}C , 97.4% of ^{15}N , and 95.6% of ^{18}O predictions fall within 1%. Compared to traditional models, OrbNet-Equi delivers about a 5 to 10-fold reduction in MAPEs, while providing

~1000-times speedup compared to DFT. These results establish OrbNet-Equi as a rapid, accurate, and robust RPFR predictor with the potential to accelerate isotopic effect modeling.

Limitations While this work demonstrates the effectiveness of OrbNet-Equi for predicting RPFRs, the model necessarily inherits the limitations of its training data, arising from both DFT and the UBM framework. First, because the QM9s-RPFR dataset is derived from DFT molecular constants, predictions remain dependent on the chosen level of theory and basis set; the implications of this dependency are discussed in a companion paper. In practice, such model-dependent biases often partially cancel when comparing RPFRs between molecules, yielding more reliable equilibrium isotope effects than absolute RPFRs [33]. Second, the UBM model relies on the Born–Oppenheimer (BO) and harmonic approximations [34]. For non-hydrogen isotopologues, corrections are generally only important at very low temperatures [34, 35], whereas for hydrogen isotopologues, neglecting anharmonic and BO corrections can introduce errors even above 200 K. These contributions, omitted in QM9s-RPFR due to their computational expense, will be evaluated in the companion paper. Nonetheless, errors from the BO and harmonic approximations tend to cancel between unsubstituted and substituted isotopologues, reducing their overall impact [36]. Addressing such limitations is an active area of research.

Possible Extensions Our dataset provides a valuable foundation for developing models to accurately predict RPFRs. However, its limited diversity of elements and the small size of its molecules present challenges, particularly since chemical reactions often involve larger species with a more diverse set of elements, such as sulfur and phosphorus. In addition, the dataset lacks charged, open-shell, and solvated species, which are commonly encountered in practical chemical reactions. These limitations can be addressed by incorporating alternative datasets that offer greater diversity in elements, larger molecules, or solvated species [37–39]. Modeling such more complex species can be done with appropriate unified models, such as OrbitAll [16]. Furthermore, including RPFRs of transition states, which are needed for modeling kinetically controlled reactions, would benefit from datasets that explicitly contain reactive species and associated transition states [40].

Furthermore, the current implementation is not applicable to predicting RPFRs of multiply substituted isotopologues. For example, predicting the RPFR of a doubly substituted isotopologue requires specifying the pair of atomic sites substituted with heavier isotopes, making it no longer a purely node-wise property. A natural extension of the current method would involve predicting an $(N_{\text{atom}})^2$ matrix, where each element represents the RPFR of the isotopologue in which the atoms corresponding to the given row and column are both substituted. This approach can be further generalized to predicting an $(N_{\text{atom}})^n$ tensor for RPFRs of n -tuply substituted isotopologues. In parallel, the framework could be extended to predict isotopic mass laws, relationships among isotope effects, that would facilitate RPFR prediction for elements with more than two stable isotopes.

5 Conclusion

We present a method for predicting RPFRs of singly substituted isotopologues by reformulating the problem as a node-wise property prediction task and employing OrbNet-Equi. The method delivers $\sim 1,000\times$ speedups over DFT overall. OrbNet-Equi achieves target accuracy with as few as 500 training molecules, and when trained on 100,000 molecules, $>95\%$ of RPFR predictions fall within experimental thresholds—representing about a 5–10 fold improvement over previous models. This orbital-based deep learning framework substantially accelerates equilibrium fractionation factor predictions, a key bottleneck for constructing equilibrium reference frames and developing mechanistic models of isotopic fingerprints unique to specific reaction networks. In doing so, our approach contributes to the work necessary for the interpretation of high-dimensional isotopic data, constraining reaction mechanisms, situating observations within equilibrium–kinetic contexts, and providing testable baselines to distinguish competing isotopic signatures, including those relevant to the search for extraterrestrial life, thereby expanding the interpretive power of isotopic measurements across geochemistry, environmental science, planetary science, and beyond.

Acknowledgements

S.A. acknowledges graduate research funding from Charlie Trimble. B.S.K. acknowledges graduate research funding from California Institute of Technology and support from the Pritzker AI+Science fund. W.A.G. acknowledges support from NSF(CBET-231117). A.A. acknowledges support from Bren endowed chair, ONR (MURI grant N00014-23-1-2654), and the Schmidt Sciences AI2050 senior fellow program.

References

- [1] Takeshi Ohba, Jun-ichi Hirabayashi, and Minoru Yoshida. “Equilibrium temperature and redox state of volcanic gas at Unzen volcano, Japan”. en. In: *Journal of Volcanology and Geothermal Research* 60.3-4 (May 1994), pp. 263–272. ISSN: 03770273. DOI: 10.1016/0377-0273(94)90055-8. URL: <https://linkinghub.elsevier.com/retrieve/pii/037702739400558> (visited on 04/16/2025).
- [2] John M. Eiler. ““Clumped-isotope” geochemistry—The study of naturally-occurring, multiply-substituted isotopologues”. en. In: *Earth and Planetary Science Letters* 262.3-4 (Oct. 2007), pp. 309–327. ISSN: 0012821X. DOI: 10.1016/j.epsl.2007.08.020. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0012821X07005109> (visited on 08/19/2021).
- [3] Elise B. Wilkes et al. “Position-specific carbon isotope analysis of serine by gas chromatography/Orbitrap mass spectrometry, and an application to plant metabolism”. en. In: *Rapid Communications in Mass Spectrometry* 36.18 (Sept. 2022), e9347. ISSN: 0951-4198, 1097-0231. DOI: 10.1002/rcm.9347. URL: <https://analyticalsciencejournals.onlinelibrary.wiley.com/doi/10.1002/rcm.9347> (visited on 05/20/2025).
- [4] Cajetan Neubauer et al. “Discovering Nature’s Fingerprints: Isotope Ratio Analysis on Bioanalytical Mass Spectrometers”. en. In: *Journal of the American Society for Mass Spectrometry* 34.4 (Apr. 2023), pp. 525–537. ISSN: 1044-0305, 1879-1123. DOI: 10.1021/jasms.2c00363. URL: <https://pubs.acs.org/doi/10.1021/jasms.2c00363> (visited on 11/17/2023).
- [5] John Eiler et al. “Analysis of molecular isotopic structures at high precision and accuracy by Orbitrap mass spectrometry”. en. In: *International Journal of Mass Spectrometry* 422 (Nov. 2017), pp. 126–142. ISSN: 13873806. DOI: 10.1016/j.ijms.2017.10.002. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1387380617303470> (visited on 03/11/2024).
- [6] John M. Eiler et al. “Frontiers of stable isotope geoscience”. en. In: *Chemical Geology* 372 (Apr. 2014), pp. 119–143. ISSN: 00092541. DOI: 10.1016/j.chemgeo.2014.02.006. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0009254114000886> (visited on 07/01/2021).
- [7] Jonathan Gropp, Mark A. Iron, and Itay Halevy. “Theoretical estimates of equilibrium carbon and hydrogen isotope effects in microbial methane production and anaerobic oxidation of methane”. en. In: *Geochimica et Cosmochimica Acta* 295 (Feb. 2021), pp. 237–264. ISSN: 00167037. DOI: 10.1016/j.gca.2020.10.018. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0016703720306505> (visited on 07/01/2021).
- [8] Yuyang He, Huiming Bao, and Yun Liu. “Predicting equilibrium intramolecular isotope distribution within a large organic molecule by the cutoff calculation”. en. In: *Geochimica et Cosmochimica Acta* 269 (Jan. 2020), pp. 292–302. ISSN: 00167037. DOI: 10.1016/j.gca.2019.10.032. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0016703719306866> (visited on 11/16/2023).
- [9] Eric M. Galimov. “CHAPTER 5 - Theoretical Model of Biological Fractionation of Isotopes”. In: *The Biological Fractionation of Isotopes*. Ed. by Eric M. Galimov. Academic Press, 1985, pp. 124–149. ISBN: 978-0-12-273970-5. DOI: <https://doi.org/10.1016/B978-0-12-273970-5.50010-4>. URL: <https://www.sciencedirect.com/science/article/pii/B9780122739705500104>.
- [10] Kristof T Schütt et al. “Schnet—a deep learning architecture for molecules and materials”. In: *The Journal of Chemical Physics* 148.24 (2018).
- [11] Johannes Gasteiger, Janek Groß, and Stephan Günnemann. “Directional Message Passing for Molecular Graphs”. In: *International Conference on Learning Representations (ICLR)*. 2020. URL: <https://openreview.net/forum?id=B1eWbxStPH>.

[12] Simon Batzner et al. “E (3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials”. In: *Nature communications* 13.1 (2022), p. 2453.

[13] Ilyes Batatia et al. “MACE: Higher order equivariant message passing neural networks for fast and accurate force fields”. In: *Advances in Neural Information Processing Systems* 35 (2022), pp. 11423–11436.

[14] Zhuoran Qiao et al. “OrbNet: Deep Learning for Quantum Chemistry Using Symmetry-Adapted Atomic-Orbital Features”. en. In: *The Journal of Chemical Physics* 153.12 (Sept. 2020). arXiv:2007.08026 [physics], p. 124111. ISSN: 0021-9606, 1089-7690. DOI: 10.1063/5.0021955. URL: <http://arxiv.org/abs/2007.08026> (visited on 09/12/2024).

[15] Zhuoran Qiao et al. “Informing geometric deep learning with electronic interactions to accelerate quantum chemistry”. en. In: *Proceedings of the National Academy of Sciences* 119.31 (Aug. 2022), e2205221119. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.2205221119. (Visited on 04/24/2024).

[16] Beom Seok Kang et al. *OrbitAll: A Unified Quantum Mechanical Representation Deep Learning Framework for All Molecular Systems*. 2025. arXiv: 2507.03853 [cs.LG]. URL: <https://arxiv.org/abs/2507.03853>.

[17] Lixue Cheng et al. “A universal density matrix functional from molecular orbital-based machine learning: Transferability across organic molecules”. In: *The Journal of chemical physics* 150.13 (2019).

[18] Fenris Lu et al. “Fast near ab initio potential energy surfaces using machine learning”. In: *The Journal of Physical Chemistry A* 126.25 (2022), pp. 4013–4024.

[19] Lixue Cheng et al. “Regression clustering for improved accuracy and training costs with molecular-orbital-based machine learning”. In: *Journal of Chemical Theory and Computation* 15.12 (2019), pp. 6668–6677.

[20] Lixue Cheng et al. “Molecular-orbital-based machine learning for open-shell and multi-reference systems with kernel addition Gaussian process regression”. en. In: *The Journal of Chemical Physics* 157.15 (Oct. 2022), p. 154105. ISSN: 0021-9606, 1089-7690. DOI: 10.1063/5.0110886. (Visited on 05/11/2024).

[21] Konstantin Karandashev and O Anatole von Lilienfeld. “An orbital-based representation for accurate quantum machine learning”. In: *The Journal of Chemical Physics* 156.11 (2022).

[22] Christian Venturella et al. “Unified Deep Learning Framework for Many-Body Quantum Chemistry via Green’s Functions”. In: *Nature Computational Science* (2025).

[23] Zihan Zou et al. “A deep learning model for predicting selected organic molecular spectra”. en. In: *Nature Computational Science* 3.11 (Nov. 2023), pp. 957–964. ISSN: 2662-8457. DOI: 10.1038/s43588-023-00550-y. URL: <https://www.nature.com/articles/s43588-023-00550-y> (visited on 03/04/2025).

[24] Raghunathan Ramakrishnan et al. “Quantum chemistry structures and properties of 134 kilo molecules”. en. In: *Scientific Data* 1.1 (Aug. 2014), p. 140022. ISSN: 2052-4463. DOI: 10.1038/sdata.2014.22. URL: <https://www.nature.com/articles/sdata201422> (visited on 09/12/2024).

[25] Stefan Grimme, Christoph Bannwarth, and Philip Shushkov. “A Robust and Accurate Tight-Binding Quantum Chemical Method for Structures, Vibrational Frequencies, and Noncovalent Interactions of Large Molecular Systems Parametrized for All spd-Block Elements (Z = 1–86)”. In: *Journal of Chemical Theory and Computation* 13.5 (2017). PMID: 28418654, pp. 1989–2009. DOI: 10.1021/acs.jctc.7b00118. URL: <https://doi.org/10.1021/acs.jctc.7b00118>.

[26] Sebastian Spicher and Stefan Grimme. “Single-point Hessian calculations for improved vibrational frequencies and rigid-rotor-harmonic-oscillator thermodynamics”. In: *Journal of Chemical Theory and Computation* 17.3 (2021), pp. 1701–1714.

[27] José Manuel Vicent-Luna, Sofia Aperi, and Shuxia Tao. “Efficient computation of structural and electronic properties of halide perovskites using density functional tight binding: GFN1-xTB method”. In: *Journal of chemical information and modeling* 61.9 (2021), pp. 4415–4424.

[28] Markus Bursch, Andreas Hansen, and Stefan Grimme. “Fast and reasonable geometry optimization of lanthanoid complexes with an extended tight binding quantum chemical method”. In: *Inorganic chemistry* 56.20 (2017), pp. 12485–12491.

[29] Maryam Nurhuda, Carole C Perry, and Matthew A Addicoat. “Performance of GFN1-xTB for periodic optimization of metal organic frameworks”. In: *Physical Chemistry Chemical Physics* 24.18 (2022), pp. 10906–10914.

[30] Markus Bursch, Hagen Neugebauer, and Stefan Grimme. “Structure optimisation of large transition-metal complexes with extended tight-binding methods”. In: *Angewandte Chemie International Edition* 58.32 (2019), pp. 11078–11087.

[31] L. Chimiak et al. “Carbon isotope evidence for the substrates and mechanisms of prebiotic synthesis in the early solar system”. en. In: *Geochimica et Cosmochimica Acta* 292 (Jan. 2021), pp. 188–202. ISSN: 00167037. DOI: 10.1016/j.gca.2020.09.026. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0016703720305901> (visited on 07/01/2021).

[32] Ross Girshick. “Fast r-cnn”. In: *Proceedings of the IEEE international conference on computer vision*. 2015, pp. 1440–1448.

[33] James R. Rustad. “Ab initio calculation of the carbon isotope signatures of amino acids”. en. In: *Organic Geochemistry* 40.6 (June 2009), pp. 720–723. ISSN: 0146-6380. DOI: 10.1016/j.orggeochem.2009.03.003. URL: <https://www.sciencedirect.com/science/article/pii/S0146638009000643> (visited on 07/01/2021).

[34] P Richet, Y Bottinga, and M Javoy. “A Review of Hydrogen, Carbon, Nitrogen, Oxygen, Sulphur, and Chlorine Stable Isotope Fractionation Among Gaseous Molecules”. en. In: *Annual Review of Earth and Planetary Sciences* 5.1 (May 1977), pp. 65–110. ISSN: 0084-6597, 1545-4495. DOI: 10.1146/annurev.ea.05.050177.000433. URL: <http://www.annualreviews.org/doi/10.1146/annurev.ea.05.050177.000433> (visited on 07/01/2021).

[35] Yining Zhang and Yun Liu. “The theory of equilibrium isotope fractionations for gaseous molecules under super-cold conditions”. en. In: *Geochimica et Cosmochimica Acta* 238 (Oct. 2018), pp. 123–149. ISSN: 00167037. DOI: 10.1016/j.gca.2018.07.001. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0016703718303776> (visited on 08/20/2025).

[36] Michael A. Webb and Thomas F. Miller. “Position-Specific and Clumped Stable Isotope Studies: Comparison of the Urey and Path-Integral Approaches for Carbon Dioxide, Nitrous Oxide, Methane, and Propane”. en. In: *The Journal of Physical Chemistry A* 118.2 (Jan. 2014), pp. 467–474. ISSN: 1089-5639, 1520-5215. DOI: 10.1021/jp411134v. URL: <https://pubs.acs.org/doi/10.1021/jp411134v> (visited on 07/01/2021).

[37] Mingzhi Yuan et al. “QMe14S: A Comprehensive and Efficient Spectral Data Set for Small Organic Molecules”. In: *The Journal of Physical Chemistry Letters* 16.16 (2025), pp. 3972–3979.

[38] Tianze Zheng et al. “Data-driven parametrization of molecular mechanics force fields for expansive chemical space coverage”. In: *Chemical Science* 16.6 (2025), pp. 2730–2740.

[39] Nicholas J Williams et al. “Hessian QM9: A quantum chemistry database of molecular Hessians in implicit solvents”. In: *Scientific data* 12.1 (2025), p. 9.

[40] Taoyong Cui et al. “Horm: A large scale molecular hessian database for optimizing reactive machine learning interatomic potentials”. In: *arXiv preprint arXiv:2505.12447* (2025).

[41] Frederick Reif. *Fundamentals of statistical and thermal physics*. Waveland Press, 2009.

[42] Sebastian Ehlert. *awvgk/tblite: Light-weight tight-binding framework*. <https://github.com/awvgk/tblite>. 2024.

[43] Frank Neese. “The ORCA program system”. In: *Wiley Interdisciplinary Reviews: Computational Molecular Science* 2.1 (2012), pp. 73–78.

A Proofs

A.1 Temperature Dependence of RPFR

According to the UBM model, the RPFR is defined by

$$\beta = \prod_{j=1}^{\alpha} \left(\frac{\omega'_j}{\omega_j} \cdot \frac{e^{-u'_j/2}}{e^{-u_j/2}} \cdot \frac{1 - e^{-u_j}}{1 - e^{-u'_j}} \right), \quad (9)$$

where $u_j = \frac{\omega_j \hbar}{k_B T}$. The vibrational temperature of molecules for the j -th vibrational mode is defined by,

$$\theta_{\text{vib},j} = \frac{\omega_j \hbar}{k_B}, \quad (10)$$

which then $u_j = \theta_{\text{vib},j}/T$. Typically, the vibrational temperatures of molecules with the relatively light atoms (H, C, N, O, F) are much larger than the room temperature [41]. This effectively makes the right side term of (9) the multiplicand to approximate to 1. I.e.,

$$\begin{aligned} \beta &= \prod_{j=1}^{\alpha} \left(\frac{\omega'_j}{\omega_j} \cdot \frac{e^{-\theta'_{\text{vib},j}/2T}}{e^{-\theta_{\text{vib},j}/2T}} \cdot \frac{1 - e^{-\theta_{\text{vib},j}/T}}{1 - e^{-\theta'_{\text{vib},j}/T}} \right), \\ &\approx \prod_{j=1}^{\alpha} \left(\frac{\omega'_j}{\omega_j} \cdot \frac{e^{-\theta'_{\text{vib},j}/2T}}{e^{-\theta_{\text{vib},j}/2T}} \right), \\ &= \prod_{j=1}^{\alpha} \left(\frac{\omega'_j}{\omega_j} \cdot \left(\frac{e^{-\theta'_{\text{vib},j}}}{e^{-\theta_{\text{vib},j}}} \right)^{1/2T} \right). \end{aligned} \quad (11)$$

Therefore,

$$\begin{aligned} \ln \beta(T) &\approx \sum_{j=1}^{\alpha} \left[\ln \left(\frac{\omega'_j}{\omega_j} \right) + \frac{\theta_{\text{vib},j} - \theta'_{\text{vib},j}}{2T} \right] \\ &\equiv \underbrace{\left(\sum_{j=1}^{\alpha} \ln \frac{\omega'_j}{\omega_j} \right)}_A + \underbrace{\frac{1}{T} \left(\frac{1}{2} \sum_{j=1}^{\alpha} (\theta_{\text{vib},j} - \theta'_{\text{vib},j}) \right)}_B. \end{aligned} \quad (12)$$

A.2 Temperature Dependence of model error

MAPE is defined by,

$$\begin{aligned} \text{MAPE}(T) &= \frac{\hat{\beta} - \beta}{\beta} = \frac{e^{\ln \hat{\beta}} - e^{\ln \beta}}{e^{\ln \beta}} \\ &= e^{\ln \hat{\beta} - \ln \beta} - 1 = e^{\Delta \ln \beta} - 1 \approx \Delta \ln \beta, \quad \text{for } |\Delta \ln \beta| \ll 1, \end{aligned} \quad (13)$$

where

$$\begin{aligned}
\Delta \ln \beta(T) &= \ln \hat{\beta}(T) - \ln \beta(T) \\
&= \left[(A + \Delta A) + \frac{B + \Delta B}{T} \right] - \left[A + \frac{B}{T} \right] \\
&= \Delta A + \frac{\Delta B}{T}.
\end{aligned} \tag{14}$$

Combining (13) and (14), we get:

$$\text{MAPE}(T) \approx \left| \Delta A + \frac{\Delta B}{T} \right|, \tag{15}$$

$$\ln(\text{MAPE}(T)) \approx \ln \left| \Delta A + \frac{\Delta B}{T} \right|, \tag{16}$$

Since $1/T$ term dominates (i.e., $\Delta A \approx 0$ or low T),

$$\ln(\text{MAPE}(T)) \approx \ln |\Delta B| - \ln T, \tag{17}$$

which shows the approximate linearity of $\ln \text{MAPE}$ linear to $-\ln T$.

As derived above, Figure 5 illustrates three key relationships: $\ln \beta$ versus $1/T$ (Equation (12)), $\ln \text{MAPE}$ versus $1/T$ (Equation (15)), and $\ln \text{MAPE}$ versus $\ln T$ (Equation (17)). Consistent with these trends, Figure 4 shows that OrbNet-Equi exhibits similar temperature dependence: the MAPE decreases approximately linearly with increasing temperature, with slopes near -1 for both the H and the non-H models.

B Data Processing

The original QM9s dataset [23] provides molecular geometries optimized at the B3LYP/def-TZVP level of theory. In this work, our objective is to predict RPFRs of different isotopologues using only their structural information—namely, atomic identities and bonding connectivity—without relying on pre-computed quantum geometries. However, OrbNet-Equi requires molecular geometries as input. To address this, we employ the GFN1-xTB framework to perform geometry optimizations, which offers an efficient and accurate approximation at a fraction of the computational cost of DFT. Prior to the GFN1-xTB optimization, we generate initial geometry guesses by re-optimizing the QM9s geometries using the universal force field (UFF) method implemented in RDKit.

All geometry optimizations are performed using the GFN1-xTB implementation in `xtb-6.7.1` [25], and orbital feature generation is carried out with `tblite-0.5.0` [42]. For the geometry optimizations and Hessian calculations at the B3LYP/def-TZVP level of theory used in the speed-up analysis, we employ `orca-6.0.0` [43].

C Training Configuration

We follow the same training configuration as described in [15], unless otherwise specified. All models are trained using a linear warm-up learning rate schedule for the first 100 epochs, followed by a cosine annealing schedule for the subsequent 400 epochs.

For the models used to generate the learning curves, the maximum learning rate is set to 0.002. In the case of models trained on RPFRs at low temperatures, we observed training instabilities due to the significantly larger magnitude of low-temperature RPFRs compared to those at high temperatures. To mitigate this divergence, we adopt smaller maximum learning rates for these models. The specific maximum learning rates used in each case are summarized in Table 2.

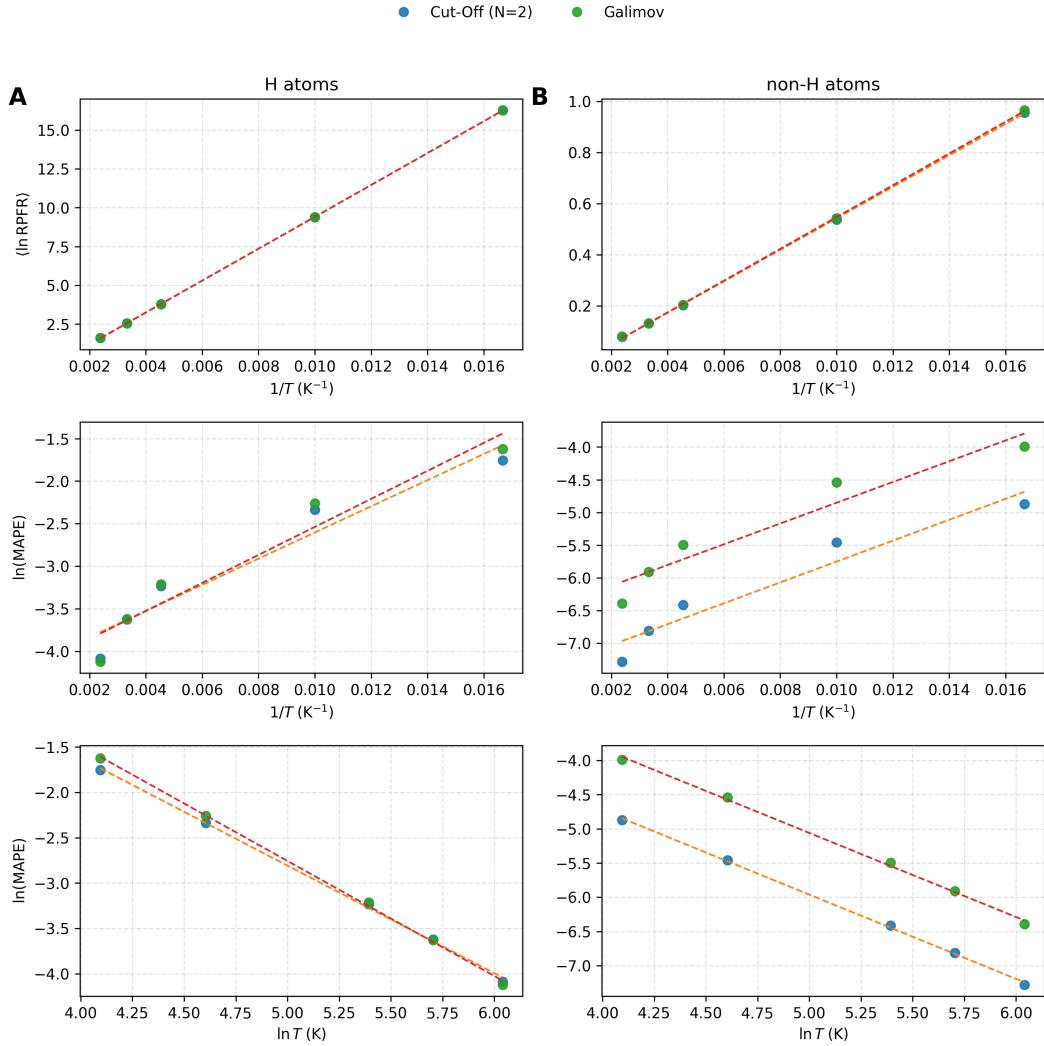


Figure 5: Temperature scaling for RPFR and mean absolute permille error. **Row 1:** Mean \ln RPFR vs $1/T$ illustrating the log-linear relationship between the inverse of temperature and RPFR following the UBM equation. **Row 2:** Mean \ln MAPE vs $1/T$ **Row 3:** $\ln(\text{MAPE})$ versus $\ln T$ (Power law). Columns split H and non-H; points are per temperature, dashed lines are OLS fits per model.

D Time Comparison

We measure the calculation times of different operations for 1,000 randomly selected molecules in the test set. For fair comparison, we start both DFT calculations and xTB-OrbNet-Equi calculations from the UFF-optimized molecular structures. The distribution of calculation times for each operation is shown in Figure 6. ‘OrbNet-Equi Total’ corresponds to the sum of ‘OrbNet-Equi Inference’ and ‘xTB Geometry Optimization’, while ‘DFT Total’ represents the sum of ‘DFT Hessian Calculation’ and ‘DFT Geometry Optimization’.

Table 2: Maximum learning rates used for each temperature (in K) RPFR predictions.

Temperature (K)	Model Type	Maximum Learning Rate	Test Set MAPE
50	H	0.0001	42.31
90	H	0.0001	11.89
130	H	0.001	7.38
170	H	0.001	5.40
210	H	0.001	4.19
250	H	0.002	3.42
290	H	0.002	2.87
330	H	0.002	2.48
10	non-H	0.001	9.88
50	non-H	0.002	1.73
90	non-H	0.002	0.92
130	non-H	0.002	0.60
170	non-H	0.002	0.44
210	non-H	0.002	0.34
250	non-H	0.002	0.28
290	non-H	0.002	0.23
330	non-H	0.002	0.19

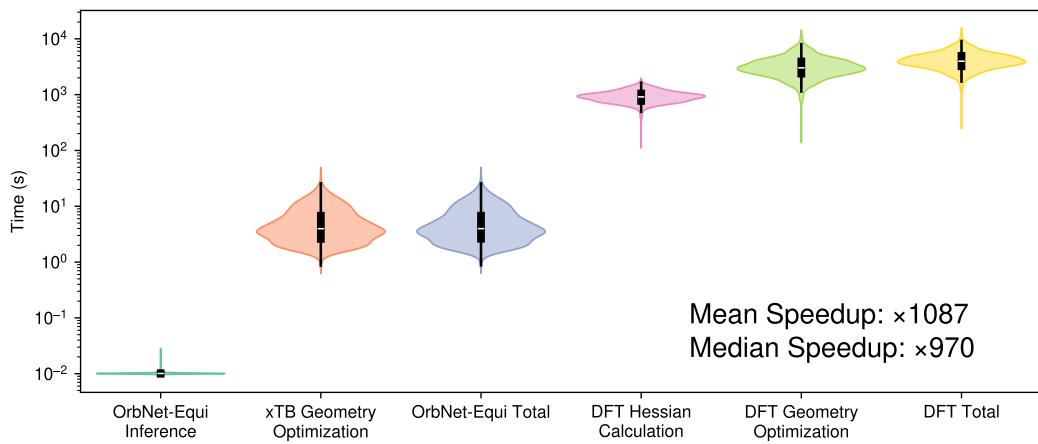


Figure 6: Distribution of calculation times for each operation.