
Hierarchical Grammar-Induced Geometry for Data-Efficient Molecular Property Prediction

Minghao Guo¹ Veronika Thost² Samuel W Song¹ Adithya Balachandran¹
Payel Das² Jie Chen² Wojciech Matusik¹

Abstract

The prediction of molecular properties is a crucial task in the field of material and drug discovery. The potential benefits of using deep learning techniques are reflected in the wealth of recent literature. Still, these techniques are faced with a common challenge in practice: Labeled data are limited by the cost of manual extraction from literature and laborious experimentation. In this work, we propose a data-efficient property predictor by utilizing a learnable hierarchical molecular grammar that can generate molecules from grammar production rules. Such a grammar induces an explicit geometry of the space of molecular graphs, which provides an informative prior on molecular structural similarity. The property prediction is performed using graph neural diffusion over the grammar-induced geometry. On both small and large datasets, our evaluation shows that this approach outperforms a wide spectrum of baselines, including supervised and pre-trained graph neural networks. We include a detailed ablation study and further analysis of our solution, showing its effectiveness in cases with extremely limited data. Code is available at <https://github.com/gmh14/Geo-DEG>.

1. Introduction

Molecular property prediction is an essential step in the discovery of novel materials and drugs, as it applies to both high-throughput screening and molecule optimization. Recent advances in machine learning, particularly deep learning, have made tremendous progress in predicting property values that are difficult to measure in reality due to the associated cost. Depending on the representation form of

molecules, various methods have been proposed, including recurrent neural networks (RNN) for SMILES strings (Lusci et al., 2013; Goh et al., 2017), feed-forward networks (FFN) for molecule fingerprints (Tao et al., 2021b;a), and, more dominantly, graph neural networks (GNN) for molecule graphs (Yang et al., 2019; Bevilacqua et al., 2022; Aldeghi & Coley, 2022; Yu & Gao, 2022). They have been employed to predict biological and mechanical properties of both polymers and drug-like molecules. Typically, these methods learn a deep neural network that maps the molecular input into an embedding space, where molecules are represented as latent features and then transformed into property values. Despite their promising performance on common benchmarks, these deep learning-based approaches require a large amount of training data in order to be effective (Audus & de Pablo, 2017; Wieder et al., 2020).

In practice, however, scientists often have small datasets at their disposal, in which case deep learning fails. Many recent studies have demonstrated that handling small dataset scenarios is a non-trivial open research problem and is of practical importance which has consistently attracted the attention of the fields of molecule and material discovery (Subramanian et al., 2016; Altae-Tran et al., 2017; Audus & de Pablo, 2017). For example, due to the difficulty of generating and acquiring data – which usually entails synthesis, wet-lab measurement, and mechanical testing – state-of-the-art works on polymer property prediction using real data are limited to only a few hundred samples (Menon et al., 2019; Chen et al., 2021). Recent deep learning research has developed various methods handling small molecular datasets, including self-supervised learning (Zhang et al., 2021; Rong et al., 2020; Wang et al., 2022; Ross et al., 2021), transfer learning (Hu et al., 2020), and few-shot learning (Guo et al., 2021b; Stanley et al., 2021). These methods involve pre-training networks on large molecular datasets before being applied to domain-specific, small-scale target datasets. However, when applied to datasets of very small size (e.g., ~ 300), most of these methods are prone to perform poorly and are statistically unstable (Hu et al., 2020). Moreover, as we will show in our experiments, these methods are less reliable when deployed on target datasets that contain significant domain gaps from the pre-trained dataset (e.g.,

¹MIT CSAIL ²MIT-IBM Watson AI Lab, IBM Research. Correspondence to: Minghao Guo <guomh2014@gmail.com>.

inconsistency in molecule sizes).

As an alternative to pure deep learning-based methods, formal grammars over molecular structures offer an explicit, explainable representation for molecules and have shown their great potential in addressing molecular tasks in a data-efficient manner (Kajino, 2019; Krenn et al., 2019; Guo et al., 2021a; 2022). A molecular grammar consists of a set of production rules that can be chained to generate molecules. The production rules, which can either be manually defined (Krenn et al., 2019; Guo et al., 2021a) or learned from data (Kajino, 2019; Guo et al., 2022), encode necessary constraints for generating valid molecular structures, such as valency restrictions. A molecular grammar has the combinatorial capacity to represent a large amount of molecules using a relatively small number of production rules. It has thus been adapted as a data-efficient generative model (Kajino, 2019; Guo et al., 2022). While molecular generation based on grammars has been widely studied, extending the data-efficiency advantage of grammars to property prediction poses a new challenge that has not yet been well-explored: Grammar-based property prediction, particularly for small datasets, requires learning how to map molecule space to potentially divergent property spaces, a task that is significantly more challenging than exploring within the molecule space as the generative model does.

Motivation. In this paper, we propose a framework for highly data-efficient property prediction based on a learnable molecular grammar. We approach this problem not by trivially deploying a separate property predictor model on top of the grammar, but instead, investigating the underlying reason why most deep-learning methods fail in data-sparse cases and identifying the key motivation for integrating the grammar into our framework: A grammar can provide an *explicit prior* on the similarity of molecular structures. For instance, two molecular structures that share a common substructure would use the same sequence of grammar production rules. As it is widely recognized in cheminformatics that molecules with similar structures have similar properties (Johnson & Maggiora, 1990; Martin et al., 2002), grammar production sequences can thus be used as a strong structural prior to predict molecular properties. We aim to develop a model that explicitly represents grammar production sequences and captures the structure-level similarity between molecules. Even from few molecules, this model is expected to reveal a wealth of information regarding the structural relationship between molecules, providing the key to a data efficient property predictor.

Framework. Figure 1 outlines our approach. At the heart of our method is a grammar-induced *geometry* (in the form of a graph) for the space of molecules. In the geometry, every path tracing from the root to a leaf represents a grammar production sequence that generates a particular molecule.

Such a geometry explicitly captures the intrinsic closeness between molecules, where structurally similar molecules are closer in distance along the geometry. In contrast to the embedding space used in most deep learning methods, our grammar-induced geometry is entirely *explicit*, which can be integrated with the property predictor by considering all molecules in the space simultaneously. However, it is infeasible to construct this geometry using existing grammars, since they incur a prohibitively high computational cost as we later demonstrate in the analysis. We propose a *hierarchical molecular grammar* consisting of two parts: a pre-defined *meta grammar* at the top and a learnable *molecular grammar* at the bottom. Our hierarchical molecular grammar, without any loss of expressiveness, is the first grammar that can be used to construct a geometry at runtime. We provide both solid theoretical and experimental evidence to demonstrate that the hierarchical molecular grammar is compact yet complete. To predict properties, we exploit the structural prior captured by the grammar-induced geometry using a graph neural diffusion model over the geometry. A joint optimization framework is proposed to learn both the geometry and the diffusion in an end-to-end manner.

Evaluation. Our evaluation covers 8 commonly used benchmark datasets, including classification and regression problems, with datasets of both small (~ 300) and large (600–1,000) sizes. Our approach significantly outperforms both competitive state-of-the-art GNN approaches and fine-tuned pretrained models. Further analysis shows that when trained on only a subset of the training data (< 100), our method can achieve performance comparable to pre-trained GNNs fine-tuned on the whole training set of the downstream prediction task, thus demonstrating the effectiveness of our method on extremely small datasets.

Contributions. **1)** Our framework, to the best of our knowledge, is the first method that leverages the data efficiency of a learnable molecular grammar for the task of molecular property prediction through learning the geometry of the space of molecules. **2)** Our hierarchical molecular grammar takes the inherent advantages of general molecular grammars, including explicitness, explanatory power, and data efficiency, but expands the utility and is more compact than all existing grammars. **3)** We show that our method achieves significantly better performance on challenging small datasets and outperforms a wide spectrum of baselines on various common benchmarks.

Related Works. Our method is mainly related to three areas: 1) machine learning, particularly graph neural networks for molecular property prediction, in which we target the same problem but propose an entirely new approach; 2) grammars for molecular machine learning, which are used in existing works for molecular generation, while we use grammars to induce a geometry for molecular property prediction; and 3)

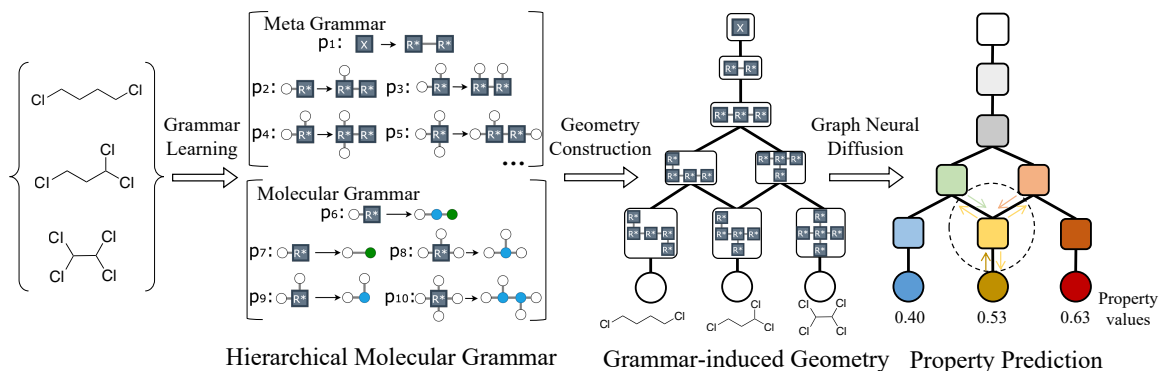


Figure 1: Overview. Given a set of molecules, we learn a hierarchical molecular grammar that can generate molecules from production rules. The hierarchical molecular grammar induces an explicit geometry for the space of molecules, where structurally similar molecules are closer in distance along the geometry. Such a grammar-induced geometry provides an informative prior for data-efficient property prediction. We achieve this by using graph neural diffusion over the geometry.

geometric deep learning, where deep networks are applied to non-Euclidean domains. We briefly discuss the latter two in Section 2 as necessary background for our approach and refer to Appendix H for an in-depth discussion.

Differences Compared to Guo et al. (2022). Our work offers key improvements over Guo et al. (2022) in multiple aspects. First, we develop a solution for property prediction, a task different from molecular generation and one that lacks a straightforward solution as to how grammars are applied. Second, we extend non-trivially a grammar to a hierarchical grammar, with crucial attributes – k -degree, edit-complete, and minimal – that provide greater expressiveness and compactness compared to the grammar in Guo et al. (2022). Third, our hierarchical grammar is generic and independent of specific molecule datasets. This allows us to construct the (meta-)geometry of the molecular graphs once in a lifetime, which is infeasible for existing, non-hierarchical grammars. Overall, our approach surpasses Guo et al. (2022) in both theory and practice, achieving compelling performance in property prediction tasks compared with a wide variety of baselines.

2. Preliminaries

Molecular Hypergraph Grammar (MHG). In MHGs, molecules are represented as hypergraphs $H = (V, E)$, where nodes represent chemical atoms and hyperedges represent chemical bonds or ring-like substructures. A MHG $G = (\mathcal{N}, \Sigma, \mathcal{P}, \mathcal{X})$ contains a set \mathcal{N} of non-terminal nodes, a set Σ of terminal nodes representing chemical atoms, and a starting node \mathcal{X} . It describes how molecular hypergraphs are generated using a set of production rules $\mathcal{P} = \{p_i | i = 1, \dots, k\}$ of form $p_i : LHS \rightarrow RHS$, where the left- and right-hand sides (LHS and RHS) are hypergraphs. Starting at \mathcal{X} , a molecule is generated by iteratively selecting a rule whose LHS matches a sub-hypergraph in

the current hypergraph and replacing it with RHS until only terminal nodes remain. For each production rule, the LHS contains only non-terminal nodes and no terminal nodes, whereas the RHS can contain both non-terminal and terminal nodes. Both sides of the production rule have the same number of anchor nodes, which indicate correspondences when LHS is replaced by RHS in a production step. For a formal definition, see (Guo et al., 2022).

Graph Diffusion models the information propagation between nodes on a graph using heat equations. The node features are updated following a diffusion PDE as follows:

$$\mathbf{U}_T = \mathbf{U}_0 + \int_0^T \frac{\partial \mathbf{U}_t}{\partial t} dt, \quad \frac{\partial \mathbf{U}_t}{\partial t} = \mathbf{L}_\alpha \mathbf{U}_t, \quad (1)$$

where matrix \mathbf{U}_t represents the features of all nodes in the graph at time t and the matrix \mathbf{L}_α , which has the same sparsity structure as a graph Laplacian, represents the diffusivity defined on all edges in the graph. \mathbf{L}_α is calculated using $a(\cdot, \cdot; \alpha)$ parameterized by α , i.e. $L_{ij} = a(U_t^{(i)}, U_t^{(j)}; \alpha)$ for all connected node pairs (i, j) . For more details, see (Chamberlain et al., 2021).

Notation. For a hypergraph grammar $G = (\mathcal{N}, \Sigma, \mathcal{P}, \mathcal{X})$, we say a graph H can be *derived* from the grammar if there is a sequence of production rules from \mathcal{P} that generates this graph, denoted by $\mathcal{X} \xrightarrow{\mathcal{P}} H$. $H_1 \xrightarrow{p} H_2$ denotes one-step grammar production that transforms graph H_1 to graph H_2 using rule p . As a special form of general graph, a tree is denoted as T and the set of all possible trees is denoted as \mathcal{T} . Note that all trees discussed in this paper are *unrooted without order*, i.e. connected acyclic undirected graphs (Bender & Williamson, 2010). $\Delta(T)$ denotes the maximal degree of T . $TED(T_1, T_2)$ denotes the *tree edit distance* between tree T_1 and tree T_2 as defined in (Zhang, 1996) and (Paaßen, 2018). $|\mathcal{P}|$ denotes the number of rules in a production rule set \mathcal{P} . $|T|$ denotes the size of a tree T .




Figure 2: (a) Example of a molecular hypergraph grammar and (b) its induced geometry. All possible production sequences to generate three molecules are outlined in the geometry, where the nodes represent intermediate or molecular hypergraphs and the edges represent production steps. Structurally similar molecules have a smaller shortest path between them and are closer in distance along the geometry.

3. Grammar-induced Geometry for Molecular Property Prediction

3.1. General Grammar-induced Geometry

Problem Formulation. A molecular property predictor can be expressed as a function $\pi(\cdot) : \mathcal{H} \rightarrow \mathbb{R}$ that maps molecules formulated as hypergraphs $H = (V, E) \in \mathcal{H}$ into scalar values. The mapping function $\pi = g \circ f$ contains two components: an embedding function $f(\cdot) : \mathcal{H} \rightarrow \mathbb{R}^n$ that maps input hypergraphs into a Euclidean *latent feature space* (also known as the *embedding space*), and a transformation function $g(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}$ (usually a simple linear function) which maps latent features into property values. The embedding function is designed to capture the similarity between molecules, such that molecules with similar properties are closer in the embedding space in terms of Euclidean distance (Cayton, 2005). The key reason why most supervised machine learning methods fail in data-sparse cases is that the embedded function fails to capture molecular similarity since it only has access to a limited number of samples with property labels. To address this issue, we propose an additional, explicit prior on molecular similarity that does not depend on data labels, but provides valuable information for property prediction.

General Grammar-induced Geometry. Our efforts are motivated by the fact that *structure-level similarity* between molecules is known to be strongly correlated with their properties, e.g., organic molecules sharing a functional group are proven to have similar chemical properties (Johnson & Maggiora, 1990; Martin et al., 2002). We *explicitly* model this structure-level similarity as an informative prior for data-efficient property prediction.

As grammars are generative models for molecules con-

structed by combining molecular substructures, we exploit the molecular hypergraph grammar to capture structural similarity between molecules. Figure 2 outlines how a molecular hypergraph grammar $G = (\mathcal{N}, \Sigma, \mathcal{P}, \mathcal{X})$ can be used to construct a *geometry* for the space of molecules in the form of a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ ¹. The geometry enumerates all possible production sequences, where every possible intermediate or molecular hypergraph is represented as a node in $\mathcal{V} = \{v | v = H_v(V_v, E_v)\}$, and each production step is represented as an edge in $\mathcal{E} = \{(s, t) | H_s \xrightarrow{p} H_t, p \in \mathcal{P}\}$. In this geometry, every leaf node represents a molecule. Every path tracing from the root $H_{root} = (\mathcal{X}, \emptyset)$ to a leaf represents a production sequence that generates the corresponding molecule. Since all the molecules can be derived from the initial node \mathcal{X} , any two leaves representing two molecules are connected by at least one common ancestor (i.e. the root) in this geometry, where the path between them represents the sequence of steps that would be required to transform from one to the other. Molecules with greater structural similarity have a smaller *shortest path* as their molecular hypergraphs share a common intermediate hypergraph. The distance between two molecules is defined as the shortest-path distance between them along the geometry with unit weight at each edge. We use this geometry as an additional input to the molecular property predictor: $\pi'(\cdot, \cdot) : \mathcal{H} \times \mathcal{G} \rightarrow \mathbb{R}$, where $\mathcal{G} \in \mathcal{G}$ is the grammar-induced geometry. The geometry can be optimized in conjunction with the property predictor in order to minimize the prediction loss. As the geometry is determined by the grammar, the optimization of the geometry can be converted into the learning of production rules, where the latter can be achieved using the method described in (Guo et al., 2022).

The crucial remaining question is how to construct this geometry $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ from a given molecular hypergraph grammar. A key characteristic of the grammar-induced geometry is that each node in the geometry represents a *unique* intermediate hypergraph or molecular hypergraph. This ensures that the geometry includes the minimal path between two molecules. To satisfy this characteristic, one trivial idea of geometry construction is to use breadth-first search (BFS). Originating from the root $H_{root} = (\mathcal{X}, \emptyset)$, the geometry can be iteratively expanded following the production rules until all the molecules involved in the property prediction task have been visited. In practice, however, such a method of constructing grammar-induced geometry is very costly, and often computationally intractable for more complex grammars. By testing with random grammars learned from (Guo et al., 2022), we find that it is infeasible to construct the geometry when there are more than ten production rules. A detailed analysis is provided in Appendix D. The ma-

¹To distinguish the graph of the geometry from the graph of molecules, we use *hypergraphs* to formulate the molecular graphs generated from grammar production.

jor bottleneck comes from the combinatorial complexity of production rules: The more rules a grammar has, the more intermediate hypergraphs it can generate. As the depth of the geometry increases, the number of intermediate hypergraphs increases exponentially. This incurs a significant computational cost, as each newly expanded hypergraph requires pair-wise isomorphism tests. This cost also poses a serious obstacle to the optimization of the geometry, since the geometry must be rebuilt at each optimization iteration.

3.2. Hierarchical Molecular Grammars

We propose a *hierarchical molecular grammar* to address the computational challenge mentioned above. Our insight stems from the fact that every hypergraph can be represented as a tree-structure object, i.e. a *junction tree* using tree decomposition. A junction tree is constructed by contracting certain vertices of the hypergraph into a single node so that it is cycle-free (Diestel, 2005). The generation process of a molecular hypergraph can thus be divided into two parts: first generating a tree with homogeneous tree nodes, then converting the tree into a molecular hypergraph by specifying a sub-hypergraph per tree node. As circle-free graphs, trees are more compact, so the possible tree structures are considerably fewer than molecular hypergraphs, and therefore can be enumerated. The key to efficient grammar-induced geometry construction is that the enumeration of all possible homogeneous tree structures is data-independent and can be computed offline. The only computational cost of constructing the geometry at run-time is attributed to converting trees into molecular hypergraphs.

We design a hierarchical molecular grammar to fulfill this principle. A hierarchical molecular grammar consists of two sub-grammars: a *meta grammar* that only generates trees, and a *molecular grammar* that can convert trees into molecular hypergraphs. The grammar-induced geometry can thus be divided into two parts: The top part (referred to as *meta geometry*) is constructed by the meta grammar, where geometry nodes represent only trees, while the bottom part (referred to as *molecular leaves*) is constructed by the molecular grammar, where geometry leaves represent molecular hypergraphs. Figure 3 shows an overview of a hierarchical molecular grammar and its induced geometry. Thanks to this hierarchical decomposition, the meta geometry can be pre-computed to apply to any molecular datasets. At run-time, we only need to construct the bottom part of the geometry, which is formulated as a byproduct of grammar learning in our approach. As a result of the geometry construction, each molecule is connected to one junction tree in the meta geometry. Molecular structure similarity is determined by the distance between their corresponding junction trees along the geometry. We provide a formal definition of meta grammar and demonstrate that despite the additional hierarchical constraint, our hierarchical molecular gram-

mar is *as expressive as* the general molecular hypergraph grammar in (Guo et al., 2022).

Meta Grammars. The definition of meta grammar is,

Definition 3.1. A *meta grammar* $\overline{G} = (\mathcal{N}, \emptyset, \mathcal{P}_{\overline{G}}, \mathcal{X})$ is a hypergraph grammar, which only contains non-terminal nodes and only generates trees, i.e. $\forall w \in \{\mathcal{X}_{\overline{G}}^{\rightarrow} w\}, w \in \mathcal{T}$.

A meta grammar \overline{G} is *k-degree* if, for all trees T of maximal degree $\Delta(T) \leq k$, we have $\mathcal{X}_{\overline{G}}^{\rightarrow} T$.

A meta grammar \overline{G} is *edit-complete* if, for any tree pair (T, T') with $|T| < |T'|$, tree edit distance $TED(T, T') = 1$, and $\mathcal{X}_{\overline{G}}^{\rightarrow} T, T'$, there is a rule $p \in \mathcal{P}_{\overline{G}}$ such that $T \xrightarrow{p} T'$.

A *k-degree, edit-complete meta grammar* \overline{G} is *minimal* if there is no other such meta grammar \overline{G}' with $|\mathcal{P}_{\overline{G}'}| < |\mathcal{P}_{\overline{G}}|$.

In Appendix I, we provide the formal construction of meta grammars for generating trees with arbitrary degree and elaborate on the three additional attributes.

Proposition 3.2. A meta grammar \overline{G} with a meta rule set $\mathcal{P}_{\overline{G}}$ as constructed in Appendix I is *k-degree (for some k), edit-complete, and minimal*.

Appendix I also provides the proof that the three attributes are satisfied. Generally speaking, we construct meta rule sets of arbitrary degrees by induction from the 1-degree meta grammar, which consists of only one meta rule. In practice, a meta grammar with a reasonable degree should be chosen to ensure the coverage of most molecules. In our experiments, we find it sufficient to use a meta rule set of degree 4, which contains 8 rules in total.

Hierarchical Molecular Grammars. We define the hierarchical molecular grammar as follows,

Proposition 3.3. A hierarchical molecular grammar $\mathbf{G} = (\overline{G}, G_{mol})$ —consisting of two hypergraph grammars: 1) a *k-degree, edit-complete, and minimal meta grammar* \overline{G} defined in Definition 3.1; and 2) a *molecular grammar* $G_{mol} = (\mathcal{N}, \Sigma, \mathcal{P}_{mol}, \emptyset)$ where the RHS of any molecular rule in \mathcal{P}_{mol} only contains terminal nodes (i.e. atoms)—is a complete grammar.

More discussion on Proposition 3.3 is in Appendix I. The completeness of our hierarchical molecular grammar also shows that there is no loss of representation capacity as compared to general molecular hypergraph grammars.

3.3. Hierarchical Grammar-induced Geometry

Geometry Construction. The geometry induced by the hierarchical molecular grammar is constructed in two phases: Figure 3(a) shows the geometry constructed by the meta grammar at the top, and the geometry constructed by the molecular grammar at the bottom. Since the meta grammar is pre-defined, the top meta geometry can be precomputed

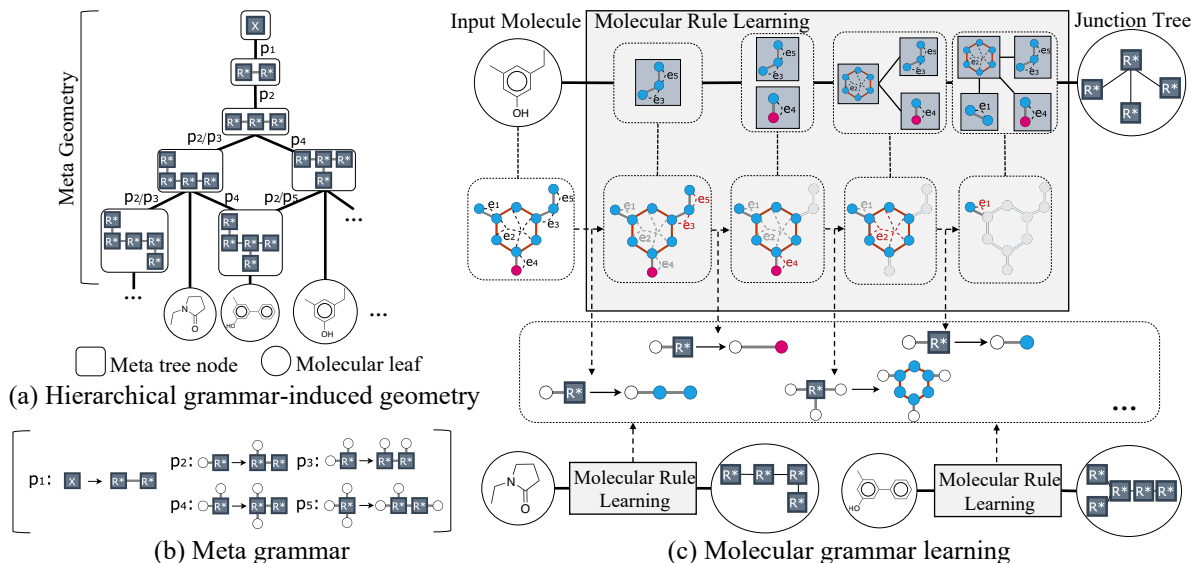


Figure 3: (a) Overview of the geometry induced by hierarchical molecular grammar. (b) A pre-defined meta grammar is used to pre-compute the meta geometry offline. (c) At run-time, molecular grammar is obtained using molecular rule learning. Each molecule is converted into a junction tree and is connected to a meta tree node of the meta geometry that represents an isomorphic meta tree.

following the depth-limited BFS procedure described in Sec. 3.1. Each node in the meta geometry represents a tree of non-terminal nodes. We call the tree generated using the meta grammar a *meta tree* and a node in the meta geometry a *meta tree node*. We find it sufficient to use a maximum BFS depth of 10 in practice. The bottom part of the geometry determines how each molecule is connected to the meta geometry as a molecular leaf through its junction tree structure. We obtain the junction trees of molecules as a result of the learning of molecular grammar rules. Thus, the geometry optimization can be achieved by learning the molecular grammar.

Molecular Grammar Learning. Figure 3(c) illustrates the process of molecular grammar learning. We follow the grammar learning from (Guo et al., 2022) but constrain the learned rules to contain only terminal nodes on the *RHS*, so as to satisfy the molecular grammar. The molecular rules are constructed along with the junction tree. Specifically, each molecule is considered as a hypergraph. At each iteration, a set of hyperedges is sampled from the hypergraph. The iterative sampling follows an i.i.d. Bernoulli distribution based on a probability function $\phi(e; \theta)$ defined on each hyperedge e with learnable parameters θ . For each connected component of these sampled hyperedges, we construct a molecular rule, where the *LHS* contains a single non-terminal node and the *RHS* contains the connected component. For the junction tree, we add a junction tree node representing the connected component and create a junction tree edge between two junction tree nodes if their corresponding connected components share a common hypergraph node.

Then the sampled hyperedges are removed from the original hypergraph. The process terminates until all hyperedges are removed. At the conclusion of the sampling process, we can obtain a junction tree of the molecular hypergraph along with a set of molecular rules. The molecule is then connected to the meta tree node in the meta geometry that represents the isomorphic meta tree to the junction tree. We provide a detailed formulation of the constructed molecular rule and the resulting junction tree in Appendix J and more details of the geometry construction in Appendix K.

3.4. Molecular Property Prediction

Graph Diffusion on Grammar-induced Geometry. Property prediction requires a model that can operate on our grammar-induced geometry and is suitable for scarce data. We choose the graph neural diffusion model GRAND (Chamberlain et al., 2021) for its effectiveness in overcoming the oversmoothing that plagues most GNNs. Three learnable components are used in a graph diffusion process: an encoder function φ defined on all the nodes in the grammar-induced geometry, a decoder function ψ defined on molecular leaves, and a graph diffusion process given in Eq. 1. Specifically, the encoder φ yields the initial state of the diffusion process $\mathbf{U}_0 = \varphi(\mathbf{U}_{in})$, where \mathbf{U}_{in} is the matrix form of the input features of all the nodes. The decoder produces predicted property values of all molecular leaves $\mathbf{u}_T = \psi(M \odot \mathbf{U}_T)$, where \mathbf{U}_T is the node-feature matrix from the final diffusion state, M is a binary mask that masks out the rows corresponding to non-molecule nodes, \odot is the Hadamard product, and \mathbf{u}_T is the resulting

vector containing property values of all molecular leaves. Our optimization objective includes the learning of both the grammar-induced geometry and the diffusion model:

$$\min_{\theta, (\varphi, \psi, \alpha)} l(\mathbf{u}_T, \hat{\mathbf{u}}) = \min_{\theta} \min_{(\varphi, \psi, \alpha)} l(\mathbf{u}_T, \hat{\mathbf{u}}), \quad (2)$$

where $\hat{\mathbf{u}}$ represents the vector of ground-truth property values for all leaves and $l(\cdot, \cdot)$ is a regression loss. Recall that the geometry is only determined by the molecular rules, so the molecular grammar learning parameters θ are the only parameters relevant for obtaining the geometry. Since θ and (φ, ψ, α) are two groups of independent variables, we exploit block coordinate descent to optimize the objective. A detailed derivation is provided in Appendix L.

Overall Framework of End-to-end Learning. First, we compute the meta geometry offline using the predefined meta grammar (Sec. 3.2). Then, given a set of molecules, we use molecular grammar learning (Sec. 3.3) to convert each molecule into a junction tree and attach it to the meta geometry to form the whole hierarchical grammar-induced geometry (Sec. 3.3). Graph diffusion (Sec. 3.4) is then applied to the geometry to obtain a scalar value for each molecular leaf, i.e. predicted property value for each molecule. Training is conducted by minimizing the error between the predicted and ground-truth values (Eq. 2). The parameters of molecular grammar learning and graph diffusion are trained in an end-to-end manner. A graphical illustration of the optimization framework is provided in Appendix L.

Application to Transductive and Inductive Setting. Our system can be used in two different ways. The first is to use our framework as both a property predictor and a generative model through the use of learned hierarchical molecular grammar, which is the default setting in our experiments. To achieve this, our framework takes all molecules as input (from both training and testing datasets), but only uses the property values of training datasets, forming a *transductive* setting. This ensures that all the input molecules can be generated using the learned grammar. The other setting is to use our framework as a pure property predictor. In this case, we ignore the learned grammar rules but consider the molecular rule learning module as a function to convert *any* input molecule (both in and out of the training dataset) into a junction tree, which is then attached to the meta geometry. Thus, our framework does not require molecules from testing datasets as input, forming an *inductive* setting. Specifically, we use the same training procedure as for the transductive case, except we only use molecules from the training dataset as input. When testing a new molecule, by using the trained grammar sampling function $\phi(e; \theta)$ and following the hyperedge contraction process, we can obtain the junction tree for that molecule. Upon attachment to the isomorphic meta tree node, the new molecule is added to grammar-induced geometry. The property value of the

molecule can be predicted following the graph diffusion using the trained decoder φ and encoder ϕ .

4. Evaluation

The experiments demonstrate the **generality of our approach** and answer the following questions: 1) Does our approach outperform existing methods on **small datasets**? 2) How well does our approach perform on **large, widely-studied datasets**? 3) To what extent is our approach effective on **extremely small datasets**? 4) Is our approach effective for **both transductive and inductive** settings?

4.1. Experiment Setup

Data. We evaluate our approach on eight datasets: CROW (a curated dataset from literature), Permeability (Yuan et al., 2021), FreeSolv (Mobley & Guthrie, 2014), Lipophilicity (Wang et al., 2015), HOPV (Lopez et al., 2016), DILI (Ma et al., 2020), PTC (Xu et al., 2018), and ClinTox (Gayvert et al., 2016). These datasets cover: 1) commonly used benchmark datasets including MolecuNet (FreeSolv, Lipophilicity, HOPV, ClinTox) and TUDataset (PTC), 2) both classification (DILI, PTC, ClinTox) and regression tasks (CROW, Permeability, FreeSolv, Lipophilicity, HOPV), and 3) sizes that are small (CROW, Permeability, HOPV, PTC, DILI) and large (FreeSolv, Lipophilicity, ClinTox). We report mean absolute error (MAE) and coefficient of determination (R^2) for regression, and Accuracy and AUC for classification. See Appendix F for details.

Baselines. We compare our approach with various approaches: Random Forest, FNN, wD-MPNN (D-MPNN), ESAN, HM-GNN, PN, and Pre-trained GIN. For descriptions, see Appendix G. To show the generality of our pipeline (called Geo-DEG), we implement two versions, each with a different diffusion encoder, GIN and MPNN. Appendix G provides the implementation details.

4.2. Results on Small Datasets

Results & Discussion. To address question 1), we conduct experiments on five small datasets: CROW, Permeability, HOPV, PTC, and DILI. Table 1 shows the results on the first four datasets and Appendix A shows the results on DILI. Both variants of our method outperform all the other methods by a large margin. For the two polymer datasets, CROW is a more challenging dataset than Permeability, as it has fewer samples and a broader range of property values (with a standard derivation of 86.8 versus 1.24 for Permeability). On CROW, traditional machine learning methods such as Random Forest and FNN are quite competitive and even outperform modern GNN architectures such as ESAN and HM-GNN. wD-MPNN achieves reasonable performance thanks to the special graph representation for polymer en-

Table 1: Results on 6 out of 8 datasets (best result **bolded**, second-best underlined). More results on the other two datasets are in Appendix A. Our approach (Geo-DEG) outperforms both state-of-the-art supervised and semi-supervised GNNs.

Datasets	Methods	Random Forest	FFN	wD-MPNN	ESAN	HM-GNN	PN (finetuned)	Pre-trained GIN (finetuned)	Geo-DEG (GIN)	Geo-DEG (MPNN)
CROW	MAE ↓	27.9 ± 3.2	24.0 ± 2.1	20.6 ± 1.3	26.1 ± 1.3	30.8 ± 1.8	21.1 ± 1.3	19.3 ± 1.4	17.0 ± 1.4	18.5 ± 1.2
	R ² ↑	0.67 ± 0.08	0.84 ± 0.02	0.89 ± 0.02	0.79 ± 0.02	0.76 ± 0.01	0.89 ± 0.01	<u>0.91 ± 0.01</u>	0.92 ± 0.01	<u>0.91 ± 0.01</u>
Permeability	MAE ↓	0.58 ± 0.01	0.56 ± 0.04	0.46 ± 0.03	0.40 ± 0.03	0.49 ± 0.03	0.48 ± 0.04	0.46 ± 0.04	<u>0.34 ± 0.02</u>	0.32 ± 0.03
	R ² ↑	0.72 ± 0.03	0.73 ± 0.06	0.80 ± 0.03	0.81 ± 0.03	0.68 ± 0.01	0.70 ± 0.03	0.69 ± 0.02	0.84 ± 0.02	<u>0.83 ± 0.02</u>
FreeSolv	MAE ↓	4.58 ± 0.34	3.67 ± 0.40	<u>0.54 ± 0.08</u>	0.67 ± 0.07	0.63 ± 0.05	0.65 ± 0.04	1.01 ± 0.11	0.62 ± 0.06	0.49 ± 0.06
	R ² ↑	0.60 ± 0.09	0.77 ± 0.05	<u>0.90 ± 0.02</u>	0.86 ± 0.03	0.86 ± 0.02	0.85 ± 0.02	0.74 ± 0.03	<u>0.90 ± 0.02</u>	0.94 ± 0.02
Lipophilicity	MAE ↓	0.64 ± 0.04	0.51 ± 0.02	<u>0.44 ± 0.02</u>	0.46 ± 0.04	0.58 ± 0.04	0.56 ± 0.04	0.52 ± 0.03	0.48 ± 0.02	0.42 ± 0.02
	R ² ↑	0.65 ± 0.03	0.80 ± 0.03	<u>0.90 ± 0.02</u>	<u>0.90 ± 0.01</u>	0.84 ± 0.02	0.81 ± 0.02	0.89 ± 0.01	0.88 ± 0.02	0.91 ± 0.02
HOPV	MAE ↓	0.36 ± 0.03	0.35 ± 0.03	0.36 ± 0.03	0.37 ± 0.02	0.40 ± 0.02	0.42 ± 0.02	0.38 ± 0.02	<u>0.32 ± 0.03</u>	0.30 ± 0.02
	R ² ↑	0.69 ± 0.05	0.67 ± 0.06	0.69 ± 0.04	0.66 ± 0.06	0.65 ± 0.05	0.65 ± 0.04	0.66 ± 0.03	<u>0.70 ± 0.03</u>	0.74 ± 0.03
PTC	Acc. ↑	0.60 ± 0.06	0.58 ± 0.06	0.67 ± 0.06	0.64 ± 0.08	<u>0.66 ± 0.07</u>	0.61 ± 0.08	0.62 ± 0.09	0.64 ± 0.09	0.69 ± 0.07
	AUC ↑	0.63 ± 0.05	0.61 ± 0.04	0.70 ± 0.05	0.68 ± 0.06	<u>0.69 ± 0.06</u>	0.65 ± 0.07	0.66 ± 0.07	0.68 ± 0.06	0.71 ± 0.07

sembles. PN and Pre-trained GIN perform exceptionally well on CROW due to their pre-training on large datasets. However, this pre-training does not help on Permeability, which has a much larger average molecule size (with an average molecular weight of 391.8 versus 153.8 for CROW). Thus, there is a domain gap between the dataset to pre-train these networks and the Permeability dataset used for testing, resulting in poor performance of both methods. ESAN benefits from the subgraph selection scheme on large molecular graphs and performs well on Permeability. The poor performance of HM-GNN on both datasets shows that it is not as effective in regression as it is in classification. The overall results show that: 1) our method can handle molecules with varying sizes, and 2) the superior performance of our method when coupled with either GIN or MPNN confirms its generalizability to different feature extractors. More discussion on the other datasets are given in Appendix A.

4.3. Results on Large Datasets

Results & Discussion. To address question 2), we conduct experiments on three large datasets: FreeSolv, Lipophilicity, and ClinTox. Table 1 and 3 show the results. Our method equipped with MPNN diffusion encoder performs the best among all approaches. On large datasets, random forest and simple FFN become less competitive than on small datasets. This is reasonable since larger datasets require larger model capacities to represent the relationship between molecules and their properties. Pre-training does not significantly help on large datasets due to the inherent domain gap between datasets used for pre-training versus those used for fine-tuning. Among the three GNN-based baselines, D-MPNN performs the best and exhibits better stability than ESAN and HM-GNN. Our approach can further improve the performance of D-MPNN and outperforms all the baselines. Results and discussion on the other datasets are given in Appendix A. These results demonstrate that our method is scalable for large molecule datasets.

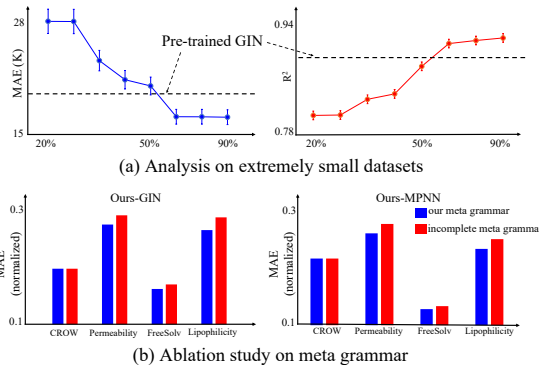


Figure 4: (a) Performance of our method trained on subsets of CROW training data with different ratios. Even when the training set is halved, our approach can still achieve competitive results compared to Pre-trained GIN *fine-tuned on the whole training set*. (b) Comparison of performance using different meta grammars, demonstrating the necessity of a complete meta grammar.

4.4. Analysis

Performance on Extremely Small Datasets. We conduct a study on the minimum viable number of samples for our method. We randomly sample 20% of CROW as a fixed testing set. Using the remaining data, we construct eight randomly selected training sets, each consisting of a portion of all remaining data, ranging from 20% to 100%. These eight training sets are used to train our method using GIN as the feature extractor. Figure 4(a) shows the performance on the testing set. Even when the training set is halved (with only 94 samples), our approach still achieves results that are comparable to those of the Pre-trained GIN *fine-tuned on the whole training set*. Appendix C includes a more detailed study for the effect of changing dataset training size on model performance.

Performance on Both Transductive and Inductive Settings. We compare both the transductive and inductive

settings of our method on the CROW dataset (the most challenging scarce dataset). Figure 5 in the Appendix shows the results. Our method achieves similar results when applied inductively and transductively, both of which outperform the best semi-supervised inductive baseline—the finetuned Pretrained GIN—on CROW. Our results show that the performance gain of our method does not depend on the transductive setting, but rather on the learnable prior introduced by the grammar-induced geometry.

Ablation Study on Meta Grammar. In this study, we examine the necessity of the meta rules in our hierarchical grammar. We remove two rules that have degree 4 on the *LHS* from the 4-degree meta grammar, resulting in a meta geometry with same number of nodes but 10% fewer edges than the one used in our main experiments. With this modified meta grammar, we run the pipeline for four datasets and compare with the original meta grammar in Figure 4(b). All four datasets exhibit a performance drop when using the modified meta grammar. The results provide experimental evidence for the necessity of a complete meta grammar.

5. Conclusion

We propose a data-efficient molecular property predictor based on a hierarchical molecular grammar. The grammar induces an explicit geometry describing the space of molecular graphs, such that a graph neural diffusion on the geometry can be used to effectively predict property values of molecules on small training datasets. One avenue of future work is to extend our pipeline to model 3D molecular structures and to address general graph design problems.

Acknowledgement

This work is supported by the MIT-IBM Watson AI Lab, and its member company, Evonik.

References

Aldeghi, M. and Coley, C. W. A graph representation of molecular ensembles for polymer property prediction. *arXiv preprint arXiv:2205.08619*, 2022.

Alon, U. and Yahav, E. On the bottleneck of graph neural networks and its practical implications. *arXiv preprint arXiv:2006.05205*, 2020.

Altae-Tran, H., Ramsundar, B., Pappu, A. S., and Pande, V. Low data drug discovery with one-shot learning. *ACS central science*, 3(4):283–293, 2017.

Audus, D. J. and de Pablo, J. J. Polymer informatics: Opportunities and challenges. *ACS macro letters*, 6(10): 1078–1082, 2017.

Bender, E. A. and Williamson, S. G. *Lists, decisions and graphs*. S. Gill Williamson, 2010.

Bevilacqua, B., Frasca, F., Lim, D., Srinivasan, B., Cai, C., Balamurugan, G., Bronstein, M. M., and Maron, H. Equivariant subgraph aggregation networks. *ICLR*, 2022.

Blakely, D., Lanchantin, J., and Qi, Y. Time and space complexity of graph convolutional networks. *Accessed on: Dec, 31, 2021*.

Bodnar, C., Di Giovanni, F., Chamberlain, B. P., Liò, P., and Bronstein, M. M. Neural sheaf diffusion: A topological perspective on heterophily and oversmoothing in gnns. *arXiv preprint arXiv:2202.04579*, 2022.

Bronstein, M. M., Bruna, J., LeCun, Y., Szlam, A., and Vandergheynst, P. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4): 18–42, 2017.

Bronstein, M. M., Bruna, J., Cohen, T., and Velicković, P. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.

Cao, W., Yan, Z., He, Z., and He, Z. A comprehensive survey on geometric deep learning. *IEEE Access*, 8:35929–35949, 2020.

Cayton, L. Algorithms for manifold learning. *Univ. of California at San Diego Tech. Rep.*, 12(1-17):1, 2005.

Chamberlain, B., Rowbottom, J., Gorinova, M. I., Bronstein, M., Webb, S., and Rossi, E. Grand: Graph neural diffusion. In *International Conference on Machine Learning*, pp. 1407–1418. PMLR, 2021.

Chen, L., Pilania, G., Batra, R., Huan, T. D., Kim, C., Kuenneth, C., and Ramprasad, R. Polymer informatics: Current status and critical next steps. *Materials Science and Engineering: R: Reports*, 144:100595, 2021.

Cordella, L. P., Foggia, P., Sansone, C., and Vento, M. An improved algorithm for matching large graphs. In *3rd IAPR-TC15 workshop on graph-based representations in pattern recognition*, pp. 149–159. Citeseer, 2001.

Cosmo, L., Kazi, A., Ahmadi, S.-A., Navab, N., and Bronstein, M. Latent-graph learning for disease prediction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 643–653. Springer, 2020.

Cruceru, C., Bécigneul, G., and Ganea, O.-E. Computationally tractable riemannian manifolds for graph embeddings. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021.

- Dai, H., Tian, Y., Dai, B., Skiena, S., and Song, L. Syntax-directed variational autoencoder for structured data. *arXiv preprint arXiv:1802.08786*, 2018.
- Diestel, R. Graph theory 3rd ed. *Graduate texts in mathematics*, 173:33, 2005.
- Elhag, A. A., Corso, G., Stärk, H., and Bronstein, M. M. Graph anisotropic diffusion for molecules. In *ICLR2022 Machine Learning for Drug Discovery*, 2022.
- Feinberg, E. N., Sur, D., Wu, Z., Husic, B. E., Mai, H., Li, Y., Sun, S., Yang, J., Ramsundar, B., and Pande, V. S. Potentialnet for molecular property prediction. *ACS central science*, 4(11):1520–1530, 2018.
- Ganea, O., Bécigneul, G., and Hofmann, T. Hyperbolic neural networks. *Advances in neural information processing systems*, 31, 2018.
- Gayvert, K. M., Madhukar, N. S., and Elemento, O. A data-driven approach to predicting successes and failures of clinical trials. *Cell chemical biology*, 23(10):1294–1301, 2016.
- Goh, G. B., Hodas, N. O., Siegel, C., and Vishnu, A. Smiles2vec: An interpretable general-purpose deep neural network for predicting chemical properties. *arXiv preprint arXiv:1712.02034*, 2017.
- Guo, M., Shou, W., Makatura, L., Erps, T., Foshey, M., and Matusik, W. Polygrammar: Grammar for digital polymer representation and generation. *arXiv preprint arXiv:2105.05278*, 2021a.
- Guo, M., Thost, V., Li, B., Das, P., Chen, J., and Matusik, W. Data-efficient graph grammar learning for molecular generation. *arXiv preprint arXiv:2203.08031*, 2022.
- Guo, Z., Zhang, C., Yu, W., Herr, J., Wiest, O., Jiang, M., and Chawla, N. V. Few-shot graph learning for molecular property prediction. In *Proceedings of the Web Conference 2021*, pp. 2559–2567, 2021b.
- Hu, W., Liu, B., Gomes, J., Zitnik, M., Liang, P., Pande, V., and Leskovec, J. Strategies for pre-training graph neural networks. *ICLR*, 2020.
- Jin, W., Barzilay, R., and Jaakkola, T. Junction tree variational autoencoder for molecular graph generation. In *International conference on machine learning*, pp. 2323–2332. PMLR, 2018.
- Jin, W., Barzilay, R., and Jaakkola, T. Hierarchical generation of molecular graphs using structural motifs. In *International conference on machine learning*, pp. 4839–4848. PMLR, 2020.
- Johnson, M. A. and Maggiora, G. M. *Concepts and applications of molecular similarity*. Wiley, 1990.
- Joo, C., Park, H., Lim, J., Cho, H., and Kim, J. Development of physical property prediction models for polypropylene composites with optimizing random forest hyperparameters. *International Journal of Intelligent Systems*, 37(6): 3625–3653, 2022.
- Kajino, H. Molecular hypergraph grammar with its application to molecular optimization. In *International Conference on Machine Learning*, pp. 3183–3191. PMLR, 2019.
- Krenn, M., Häse, F., Nigam, A., Friederich, P., and Aspuru-Guzik, A. Selfies: a robust representation of semantically constrained graphs with an example application in chemistry. *arXiv preprint arXiv:1905.13741*, 2019.
- Lopez, S. A., Pyzer-Knapp, E. O., Simm, G. N., Lutzow, T., Li, K., Seress, L. R., Hachmann, J., and Aspuru-Guzik, A. The harvard organic photovoltaic dataset. *Scientific data*, 3(1):1–7, 2016.
- Lusci, A., Pollastri, G., and Baldi, P. Deep architectures and deep learning in chemoinformatics: the prediction of aqueous solubility for drug-like molecules. *Journal of chemical information and modeling*, 53(7):1563–1575, 2013.
- Ma, H., An, W., Wang, Y., Sun, H., Huang, R., and Huang, J. Deep graph learning with property augmentation for predicting drug-induced liver injury. *Chemical Research in Toxicology*, 34(2):495–506, 2020.
- Martin, Y. C., Kofron, J. L., and Traphagen, L. M. Do structurally similar molecules have similar biological activity? *Journal of medicinal chemistry*, 45(19):4350–4358, 2002.
- Maziarz, K., Jackson-Flux, H., Cameron, P., Sirockin, F., Schneider, N., Stiefl, N., Segler, M., and Brockschmidt, M. Learning to extend molecular scaffolds with structural motifs. *arXiv preprint arXiv:2103.03864*, 2021.
- Menon, A., Thompson-Colón, J. A., and Washburn, N. R. Hierarchical machine learning model for mechanical property predictions of polyurethane elastomers from small datasets. *Frontiers in Materials*, 6:87, 2019.
- Mobley, D. L. and Guthrie, J. P. Freesolv: a database of experimental and calculated hydration free energies, with input files. *Journal of computer-aided molecular design*, 28(7):711–720, 2014.
- Nigam, A., Pollice, R., Krenn, M., dos Passos Gomes, G., and Aspuru-Guzik, A. Beyond generative models: superfast traversal, optimization, novelty, exploration and discovery (stoned) algorithm for molecules using selfies. *Chemical science*, 2021.

- Paaßen, B. Revisiting the tree edit distance and its backtracing: A tutorial. *arXiv preprint arXiv:1805.06869*, 2018.
- Rong, Y., Bian, Y., Xu, T., Xie, W., Wei, Y., Huang, W., and Huang, J. Self-supervised graph transformer on large-scale molecular data. *Advances in Neural Information Processing Systems*, 33:12559–12571, 2020.
- Ross, J., Belgodere, B., Chenthamarakshan, V., Padhi, I., Mroueh, Y., and Das, P. Do large scale molecular language representations capture important structural information? *arXiv preprint arXiv:2106.09553*, 2021.
- Shervashidze, N., Schweitzer, P., Van Leeuwen, E. J., Mehlhorn, K., and Borgwardt, K. M. Weisfeiler-lehman graph kernels. *Journal of Machine Learning Research*, 12(9), 2011.
- Stanley, M., Bronskill, J. F., Maziarz, K., Misztela, H., Lanini, J., Segler, M., Schneider, N., and Brockschmidt, M. Fs-mol: A few-shot learning dataset of molecules. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- Subramanian, G., Ramsundar, B., Pande, V., and Denny, R. A. Computational modeling of β -secretase 1 (bace-1) inhibitors using ligand based approaches. *Journal of chemical information and modeling*, 56(10):1936–1949, 2016.
- Tao, L., Chen, G., and Li, Y. Machine learning discovery of high-temperature polymers. *Patterns*, 2(4):100225, 2021a.
- Tao, L., Varshney, V., and Li, Y. Benchmarking machine learning models for polymer informatics: An example of glass transition temperature. *Journal of Chemical Information and Modeling*, 61(11):5395–5413, 2021b.
- Wang, J.-B., Cao, D.-S., Zhu, M.-F., Yun, Y.-H., Xiao, N., and Liang, Y.-Z. In silico evaluation of logd7. 4 and comparison with other prediction methods. *Journal of Chemometrics*, 29(7):389–398, 2015.
- Wang, Y., Wang, J., Cao, Z., and Barati Farimani, A. Molecular contrastive learning of representations via graph neural networks. *Nature Machine Intelligence*, 4(3):279–287, 2022.
- Wieder, O., Kohlbacher, S., Kuenemann, M., Garon, A., Ducrot, P., Seidel, T., and Langer, T. A compact review of molecular property prediction with graph neural networks. *Drug Discovery Today: Technologies*, 37:1–12, 2020.
- Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256, 1992.
- Wu, F., Zhang, Q., Radev, D., Cui, J., Zhang, W., Xing, H., Zhang, N., and Chen, H. Molformer: Motif-based transformer on 3d heterogeneous molecular graphs. *arXiv preprint arXiv:2110.01191*, 2021.
- Xu, C., Liu, Q., Huang, M., and Jiang, T. Reinforced molecular optimization with neighborhood-controlled grammars. *Advances in Neural Information Processing Systems*, 33:8366–8377, 2020.
- Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*, 2018.
- Yang, K., Swanson, K., Jin, W., Coley, C., Eiden, P., Gao, H., Guzman-Perez, A., Hopper, T., Kelley, B., Mathea, M., et al. Analyzing learned molecular representations for property prediction. *Journal of chemical information and modeling*, 59(8):3370–3388, 2019.
- Yu, Z. and Gao, H. Molecular representation learning via heterogeneous motif graph neural networks. In *International Conference on Machine Learning*, pp. 25581–25594. PMLR, 2022.
- Yuan, Q., Longo, M., Thornton, A. W., McKeown, N. B., Comesana-Gandara, B., Jansen, J. C., and Jelfs, K. E. Imputation of missing gas permeability data for polymer membranes using machine learning. *Journal of Membrane Science*, 627:119207, 2021.
- Zhang, K. A constrained edit distance between unordered labeled trees. *Algorithmica*, 15(3):205–222, 1996.
- Zhang, Z., Liu, Q., Wang, H., Lu, C., and Lee, C.-K. Motif-based graph self-supervised learning for molecular property prediction. *Advances in Neural Information Processing Systems*, 34:15870–15882, 2021.
- Zhou, G., Gao, Z., Ding, Q., Zheng, H., Xu, H., Wei, Z., Zhang, L., and Ke, G. Uni-mol: A universal 3d molecular representation learning framework. 2022.

A. Additional Results

Table 2: Results on DILI.

	DILI	
	Accuracy \uparrow	ROC-AUC \uparrow
Random Forest	0.70 \pm 0.09	0.80 \pm 0.06
FFN	0.69 \pm 0.07	0.78 \pm 0.07
D-MPNN	0.75 \pm 0.10	0.83 \pm 0.07
ESAN	0.74 \pm 0.10	0.82 \pm 0.10
HM-GNN	<u>0.76 \pm 0.08</u>	0.83 \pm 0.09
PN (finetued)	0.75 \pm 0.10	0.83 \pm 0.06
Pre-trained GIN (finetued)	0.74 \pm 0.09	0.82 \pm 0.06
Geo-DEG (GIN)	<u>0.76 \pm 0.09</u>	<u>0.84 \pm 0.05</u>
Geo-DEG (MPNN)	0.78 \pm 0.08	0.86 \pm 0.06

Table 3: Results on ClinTox.

Methods	ROC-AUC \uparrow
D-MPNN	90.6 \pm 0.6
AttentiveFP	84.7 \pm 0.3
N-Gram _{RF}	77.5 \pm 4.0
N-Gram _{XGB}	87.5 \pm 2.7
GROVER _{base}	81.2 \pm 3.0
GROVER _{large}	76.2 \pm 3.7
GraphMVP	79.1 \pm 2.8
MolCLR	91.2 \pm 3.5
GEM	90.1 \pm 1.3
Pre-trained GIN	72.6 \pm 1.5
Uni-Mol	<u>91.9 \pm 1.8</u>
Geo-DEG (GIN)	74.4 \pm 1.8
Geo-DEG (MPNN)	92.2 \pm 0.7

Metrics/Methods	MAE \downarrow	R ² \uparrow
Pre-trained GIN (fine-tuned)	19.3 \pm 1.4	0.91 \pm 0.01
Geo-DEG GIN (transductive)	17.0 \pm 1.4	0.92 \pm 0.01
Geo-DEG GIN (inductive)	17.2 \pm 1.3	0.92 \pm 0.01




Figure 5: Comparison of transductive and inductive settings (left) and analysis on combining with generative models (right). Two objectives from the Permeability dataset are predicted using our method combined with the grammar-based generative model. Molecules lying on the Pareto front are highlighted.

Results & Discussion. For HOPV, since it is a small regression dataset, traditional methods (random forest and FFN) achieve competitive performance and are even better than several graph neural networks including ESAN and HM-GNN. Both variants of our method outperform other baselines by a large margin, which aligns with the results in the other two small regression datasets. In the two small classification datasets, DILI and PTC, there is no significant difference between baseline models regarding performance. The reason for this is that the test dataset is small, and the accuracy of the model can be limited by a few hard examples that cannot be classified correctly. It is therefore more informative to use the ROC-AUC score. Among the baselines, HM-GNN performs the best due to its motif-based representation. Our method can still outperform all other methods regarding both accuracy and ROC-AUC score. For ClinTox in Table 3, we compare our approach with a wide range of highly performant baseline methods from a recent paper (Zhou et al., 2022), which by itself proposes one of the state-of-the-art methods on ClinTox. Note that these are very strong baseline methods, including many methods pre-trained using 3D molecular data. The overall results show that our method achieves significantly better performance on challenging small regression datasets and outperforms a wide spectrum of baselines on various common benchmarks.

B. Combination with Generative Models.

In our optimization framework, grammar metrics can be considered as additional objectives, allowing us to jointly optimize both generative models and property predictors. Following (Guo et al., 2022), we use diversity and Retro* score as our grammar metrics and then perform joint optimization on the Permeability dataset. After training, we generate 400 molecules and predict their property values using our approach, including all six property types from Permeability. Figure 5(b) shows two out of six predicted property values with the Pareto front highlighted. Clearly, our approach can be combined with generative models to provide a comprehensive pipeline for the discovery of optimal molecules. The use of Retro* also enables finding synthesis paths of generated molecules as shown in Appendix M.

C. Analysis on Training Dataset Size

We further extend our study in Section 4.4 for the effect of changing dataset training size on model performance by including the analysis of all three small regression datasets with more baselines compared. Figure 6 illustrates the results of Pre-trained GIN, wD-MPNN/D-MPNN, and two variants of our model on CROW, Permeability, and HOPV. We report the performance of each model trained using different numbers of training samples, randomly sampled from the original training dataset. Across different training dataset sizes, our proposed method consistently outperforms other baselines. A smaller training dataset leads to a larger performance improvement, demonstrating the data-efficiency of our model.




Figure 6: Analysis on training dataset size.

D. Analysis on Construction Cost Reduction of Hierarchical Molecular Grammar

In this section, we provide empirical evidence on the construction cost reduction of our hierarchical molecular grammar compared with non-hierarchical molecular grammar from (Guo et al., 2022). We conduct ten groups of experiments by randomly sampling five molecules from CROW dataset by ten times. For each group of experiment, we sample ten grammars for both hierarchical and non-hierarchical versions respectively and construct the geometry to cover the five molecules. The non-hierarchical grammar is sampled using the algorithm from (Guo et al., 2022). The geometry construction is achieved following the BFS procedure described in Section 3.1 for both kinds of grammar. In the experiments, we find that for non-hierarchical grammar, even for five molecules, it occurs frequently that the size of the geometry grows extremely large but still cannot cover the five molecules. Therefore, we stop the construction when the size of the geometry reaches 2×10^5 . Figure 7 illustrates (a) the average running time, (b) the average number of production rules, and (c) the average number of intermediate graphs of the final constructed geometry. It can be noted that our hierarchical molecular grammar greatly reduces the running time of the geometry construction. For non-hierarchical grammars, it is always intractable to construct the geometry since the number of intermediate graphs is enormous (some approaching 1.5×10^5). Our hierarchical molecular grammar provides a practical framework for grammar-induced geometry with theoretical foundations.

E. Analysis on Computational Complexity

The computation of our approach consists of three parts: 1) the construction of grammar-induced geometry, 2) diffusion encoder, and 3) graph diffusion over the geometry. Following (Blakely et al., 2021), suppose we have K samples of molecules, each with N nodes and $|\mathcal{E}|$ edges. To construct the grammar-induced geometry, we need to sample all the




Figure 7: Analysis on construction cost reduction of hierarchical molecular grammar. For our hierarchical molecular grammar, the number of intermediate graphs is the same as the number of molecule samples, which is five in the experiments.

edges in order to contract molecules into junction trees, so the computational complexity is $\mathcal{O}(K|\mathcal{E}|)$. For diffusion encoder, the complexity is the same with general graph neural networks, which is $\mathcal{O}(KLd(Nd + |\mathcal{E}|))$, where L is the number of layers and d is the feature dimension. As indicated in (Chamberlain et al., 2021), the complexity of graph diffusion is $\mathcal{O}(|\mathcal{E}'|d(E_b + E_f))$, where \mathcal{E}' is the edge set of grammar-induced geometry, and E_b, E_f are the numbers of function evaluations for forward and backward pass, respectively. Therefore, the overall computational cost of our method is $\mathcal{O}(K|\mathcal{E}| + KLd(Nd + |\mathcal{E}|) + |\mathcal{E}'|d(E_b + E_f)) = \mathcal{O}(KLd(Nd + |\mathcal{E}|) + |\mathcal{E}'|d(E_b + E_f))$. Compared with general graph neural networks which have a computational complexity of $\mathcal{O}(KLd(Nd + |\mathcal{E}|))$, our approach entails additional computations $\mathcal{O}(|\mathcal{E}'|d(E_b + E_f))$, where $|\mathcal{E}'|$ is the sum of the number of edges in the meta geometry $|\mathcal{E}_{meta}|$ and the number of samples K , since each molecule is connected to the meta geometry by an edge. In our experiments, we use a 4-degree meta grammar with depth 10, which gives us a meta geometry of 455 edges, i.e. $455 = |\mathcal{E}_{meta}| \ll K|\mathcal{E}|$. We also find in our experiments that $E_b + E_f \ll LNd$. As a result, the overall computational complexity of our method is the same as that of general graph neural networks in the big \mathcal{O} sense. In practice, geometry can be constructed in parallel, which reduces the running time of our entire method.

F. Details on Datasets

The Chemical Retrieval on the Web (CROW) polymer database² curates a database of thermo-physical data for over 250 different polymers, focusing primarily on the most commonly used plastics and resins in industry. CROW distinguishes from other digital polymer property databases due to its heavy usage of real, experimental data individually sourced via extensive literature search. In the CROW polymer database, approximately 95% of reported polymers offer some experimental data from literature, with the remaining 5% derived purely from simulation. We only use those data with labels from real experiments in our evaluation. The CROW polymer database reports several properties including the glass transition temperature, Hildebrand solubility parameter, molar heat capacity, refractive index, and molar cohesive energy. Of these, we choose glass transition temperature as our benchmark performance comparison since it is one of the best documented polymer properties in the literature (Tao et al., 2021b). For each dataset except ClinTox, we randomly split the data into 4 : 1 training and testing sets and create five such splits using 5 random seeds. For ClinTox, we follow the settings in (Zhou et al., 2022) to train our model and report the results of other baselines directly from the paper. Each of five splits is used to train and test a separate instance of each benchmark model. For the other three datasets, we refer the reader to the original papers: Permeability (Yuan et al., 2021), FreeSolv (Mobley & Guthrie, 2014), Lipophilicity (Wang et al., 2015), HOPV (Lopez et al., 2016), DILI (Ma et al., 2020), and PTC (Xu et al., 2018).

²<https://polymerdatabase.com/>

G. Details on the Implementation

Baselines. We compare our approach with various baselines: 1) Random Forest and FFN, two of the best-performing machine learning models for polymer informatics benchmarking (Tao et al., 2021b); 2) wD-MPNN, a state-of-the-art method specifically designed for polymer property prediction (Aldeghi & Coley, 2022); 3) ESAN, a general GNN architecture with enhanced expressive power (Bevilacqua et al., 2022); 4) HM-GNN, a motif-based GNN for molecular feature representation (Yu & Gao, 2022); 5) PN and Pre-trained GIN, two pre-trained GNNs with state-of-the-art performance on few-shot learning (Stanley et al., 2021) and transfer learning (Hu et al., 2020). As HM-GNN and Pre-trained GIN only provide code for graph classification, we modify the final layer of their networks and use l_1 loss for training. For the other methods, we follow the same implementation as their original papers. Since wD-MPNN cannot be deployed on general molecules other than polymers, we report the results of D-MPNN (Yang et al., 2019) instead for two large datasets.

Our System. For our approach, we use 4-degree meta grammar, which contains eight rules. The meta geometry contains all the meta trees whose size is smaller than 11, resulting in 149 nodes and 455 edges. For the molecular rule learning of θ , we follow all the hyperparameters used in (Guo et al., 2022). For the graph diffusion, the input feature of each meta tree node is the Weisfeiler Lehman graph hashing feature (Shervashidze et al., 2011). The encoder for meta tree nodes is an embedding layer that maps hashing features into a 300-dimension continuous vector. For molecular leaves, we use two different encoders: GIN from (Xu et al., 2018) and MPNN from (Yang et al., 2019), both of which output a feature vector of dimension 300. For the decoder, we use a one-layer fully connected network with size 300. For the graph diffusion process, we follow (Chamberlain et al., 2021) and use Dormand–Prince adaptive step size scheme (DIORIS) with adjoint method. The diffusivity function $a(\cdot, \cdot; \alpha)$ is an attention function. We use Adam optimizer for the training of θ and (φ, ψ, α) , with learning rate 0.01 and 0.001, respectively. We train θ for ten epochs. For each training epoch of θ , we train (φ, ψ, α) for 50 epochs.

H. Related works

Machine Learning for Molecular Property Prediction. The use of machine learning methods to predict molecular properties has a long history. Many early methods use SMILES strings or handcrafted fingerprints as input and rely on traditional machine learning methods, such as random forest and Gaussian processes, which are still competitive today in many applications (Tao et al., 2021b; Joo et al., 2022). Recently, graph-based representations of molecules have gained increasing popularity with the development of GNNs (Feinberg et al., 2018; Xu et al., 2018; Wu et al., 2021; Bevilacqua et al., 2022; Yu & Gao, 2022; Aldeghi & Coley, 2022; Alon & Yahav, 2020). We refer the reader to (Wieder et al., 2020) for a detailed review of GNN-based property predictors. Current state-of-the-art GNN-based methods provide more advanced representations of molecules based on graphs. (Bevilacqua et al., 2022) represents each individual graph as a set of subgraphs, increasing the expressive power of GNNs. (Aldeghi & Coley, 2022) tailors molecular graphs by adding stochastic edges and designs GNN specifically for polymers based on (Yang et al., 2019). (Yu & Gao, 2022) and (Zhang et al., 2021) leverage motifs to model molecular relationships. All these methods require large training datasets to achieve reasonable performance. It is common to use self-supervised learning (Rong et al., 2020; Wang et al., 2022) and transfer learning (Hu et al., 2020) to deal with sparse data, where neural networks are pre-trained on large datasets and then fine-tuned on the target small dataset. Most of these methods, however, focus solely on molecular classification, while our approach can also address regression problems with extremely sparse data, which are considerably more challenging.

Molecular Grammars. As an interpretable and compact design model, grammar has recently gained increasing attention in the field of molecule discovery (Dai et al., 2018; Kajino, 2019; Krenn et al., 2019; Xu et al., 2020; Nigam et al., 2021; Guo et al., 2021a; 2022). A molecular grammar can be considered as a generative model that uses production rules to generate molecules. The rules of a grammar can be either manually constructed (Dai et al., 2018; Krenn et al., 2019; Guo et al., 2021a) or automatically learned from molecular datasets (Kajino, 2019; Xu et al., 2020; Nigam et al., 2021; Guo et al., 2022). Recent works have demonstrated that a learnable hypergraph grammar can be effective in capturing hard chemical constraints, such as valency restrictions (Kajino, 2019; Guo et al., 2022). (Guo et al., 2022) further proposes a learning pipeline for constructing a hypergraph grammar-based generative model that can incorporate domain-specific knowledge from very small datasets with dozens of samples. Despite the inherent advantages of molecular grammars, such as explicitness, explanatory power, and data efficiency, existing works mainly use them for molecular generation. Prediction and optimization of molecular properties can only be accomplished by using a separate model based on the grammar representation in a latent space, which is learned individually (Kajino, 2019; Xu et al., 2020). We integrate molecular grammar into property prediction tasks by constructing a geometry of molecular graphs based on the learnable grammar,

which allows us to optimize both the generative model and the property predictor simultaneously. On extremely small datasets, our approach benefits especially from the data efficiency of grammar and achieves superior performance over existing property predictors.

Hierarchical Molecular Generation. The decomposition of molecular structures in our hierarchical molecular grammar is related to substructure-based methods in molecular generation (Jin et al., 2018; 2020; Maziarz et al., 2021). (Jin et al., 2018) employs encoders and decoders to generate a junction tree-structured scaffold and then combine chemical substructures into a molecule. (Jin et al., 2020) uses a hierarchical graph encoder-decoder to generate molecules in a coarse-to-fine manner, from atoms to connected motifs. (Maziarz et al., 2021) integrates molecule fragments and atom-by-atom construction to generate new molecules. Our approach is fundamentally different from all these existing hierarchical representations in two respects: 1) *Without* the need for any training, our proposed meta grammar (the coarse level) can enumerate all possible junction tree structures by using a compact set of meta production rules, which can be *theoretically guaranteed*. Our method only requires learning at the fine level, which are molecular fragments determined by molecular rules, whereas existing methods require learning two models for both coarse and fine levels. 2) As opposed to existing methods that use latent spaces, our meta geometry induced by meta grammar explicitly models the similarity between molecules by using graph distance along the geometry. Due to the edit-completeness of meta grammar in Definition 3.1, the graph distance is an explainable measurement of minimal editing distance between molecular graphs, whereas the distance in latent spaces used in existing methods lacks explainability.

Geometric Deep Learning applies deep neural networks to non-Euclidean domains such as graphs and manifolds with a wide range of applications (Bronstein et al., 2017; Cao et al., 2020; Bronstein et al., 2021). Related to our method, a series of recent works use graph neural diffusion by treating GNNs as a discretisation of an underlying heat diffusion PDE (Chamberlain et al., 2021; Elhag et al., 2022; Bodnar et al., 2022). These methods work on large graph data where the graph connectivity of individual nodes is provided. Additionally, there are also existing works on inferring the underlying geometry of graph data. (Ganea et al., 2018) and (Cruceru et al., 2021) embed graphs into non-Euclidean manifolds such as hyperbolic and elliptical spaces and optimize their embeddings within these Riemannian spaces. (Cosmo et al., 2020) learns a latent graph to model the underlying relationship between data samples and applies GNN to it. Different from these methods, our approach uses explicit and learnable intrinsic geometry based on graph grammar to model the relationship between molecular data. Since graph grammar is a generative model, our framework is capable of optimizing both molecular generation and property prediction simultaneously.

I. More Details of Hierarchical Molecular Grammar

Discussion on Hierarchical Molecular Grammar. The three additional attributes in Definition 3.1 can be used to deduce several desirable properties of a meta grammar. "Degree k " ensures the meta grammar is expressive and complete, covering all possible trees under a simple tree degree constraint. "Edit completeness" enables the explicit capture of transformations between two trees with edit distance one and therefore between two arbitrary trees with arbitrary distances. "Minimality" guarantees that the meta grammar is compact. Based on these three attributes, we can construct a generic meta grammar generating trees of non-terminal nodes by using only a small but expressive set of production rules. Figure 3(b) shows a 3-degree meta grammar. Each rule has one non-terminal node \mathcal{R}^* on the *LHS* and two \mathcal{R}^* s on the *RHS*. Depending on the number of anchor nodes on the *LHS*, these rules can transform tree nodes of different degrees by attaching a new node using different schemes indicated by the *RHS*. It is evident that we can generate any possible tree with a degree smaller than 4 by adopting a sequence of rules from this 3-degree meta grammar. In Proposition 3.3, "Completeness" states that any molecular graph can be derived from a hierarchical molecular grammar with a set of appropriate molecular rules. This can be demonstrated by the fact that we can generate an arbitrary molecule by first using the meta grammar to generate a tree of non-terminal nodes (which has the same tree structure as a junction tree decomposed from the molecule), and then using the molecular grammar to transform the tree into a molecular hypergraph by specifying a molecular fragment for each non-terminal node.

Meta Rules Construction. We refer to the production rule set of k -degree, edit-complete, and minimal meta grammar as a k -degree meta rule set. We visualize 4-degree meta rules in Figure 8a, where $\{p_1\}$, $\{p_1, p_2, p_3\}$, and $\{p_1, p_2, p_3, p_4, p_5\}$ correspond to meta rule sets of degree 1 to 3, respectively. 1-degree meta rule set $\mathcal{P}_{\overline{G}}^{(1)}$ is constructed as

$$\begin{aligned} \mathcal{P}_{\overline{G}}^{(1)} &= \mathcal{P}^{(1)} = \{p_0^{(1)}\}, \\ p_0^{(1)} : LHS_0^{(1)} &\rightarrow RHS_0^{(1)}, LHS_0^{(1)} := (\{\mathcal{X}\}, \emptyset), RHS_0^{(1)} := H(V_{R,0}^{(1)}, E_{R,0}^{(1)}), \\ V_{R,0}^{(1)} &= \{\mathcal{R}_1^*, \mathcal{R}_2^*\}, E_{R,0}^{(1)} = \{(\mathcal{R}_1^*, \mathcal{R}_2^*)\}. \end{aligned}$$

When using the meta grammar for production, we treat $\mathcal{R}_1^*, \mathcal{R}_2^*$ as the same type of non-terminal node, i.e. $\mathcal{R}^* = \mathcal{R}_i^*, i = 1, 2$, despite the use of indices in the above formulation.

For k -degree meta rule set $\mathcal{P}_{\overline{G}}^{(k)}$ ($k > 1$), the construction is achieved by induction:

$$\begin{aligned} \mathcal{P}_{\overline{G}}^{(k)} &= \mathcal{P}_{\overline{G}}^{(k-1)} \cup \mathcal{P}^{(k)}, \mathcal{P}^{(k)} = \{p_0^{(k)}\} \cup \bigcup_{i=1}^{\lfloor \frac{k}{2} \rfloor} \{p_i^{(k)}\}, \\ p_0^{(k)} : LHS_0^{(k)} &\rightarrow RHS_0^{(k)}, LHS_0^{(k)} := H(V_{L,0}^{(k)}, E_{L,0}^{(k)}), RHS_0^{(k)} := H(V_{R,0}^{(k)}, E_{R,0}^{(k)}), \\ p_i^{(k)} : LHS_i^{(k)} &\rightarrow RHS_i^{(k)}, LHS_i^{(k)} := H(V_{L,i}^{(k)}, E_{L,i}^{(k)}), RHS_i^{(k)} := H(V_{R,i}^{(k)}, E_{R,i}^{(k)}), i = 1, \dots, \lfloor \frac{k}{2} \rfloor, \\ V_{L,0}^{(k)} &= \{\mathcal{R}^*\} \cup \bigcup_{j=1}^{k-1} \{V_{anc,j}\}, E_{L,0}^{(k)} = \bigcup_{j=1}^{k-1} \{(\mathcal{R}^*, V_{anc,j})\}, \\ V_{R,0}^{(k)} &= \{\mathcal{R}_1^*, \mathcal{R}_2^*\} \cup \bigcup_{j=1}^{k-1} \{V_{anc,j}\}, E_{R,0}^{(k)} = \{(\mathcal{R}_1^*, \mathcal{R}_2^*)\} \cup \bigcup_{j=1}^{k-1} \{(V_{anc,j}, \mathcal{R}_2^*)\}, \\ V_{L,i}^{(k)} &= \{\mathcal{R}^*\} \cup \bigcup_{j=1}^k \{V_{anc,j}\}, E_{L,i}^{(k)} = \bigcup_{j=1}^k \{(\mathcal{R}^*, V_{anc,j})\}, \\ V_{R,i}^{(k)} &= \{\mathcal{R}_1^*, \mathcal{R}_2^*\} \cup \bigcup_{j=1}^k \{V_{anc,j}\}, E_{R,i}^{(k)} = \{(\mathcal{R}_1^*, \mathcal{R}_2^*)\} \cup \bigcup_{j=1}^i \{(V_{anc,j}, \mathcal{R}_1^*)\} \cup \bigcup_{j=i+1}^k \{(V_{anc,j}, \mathcal{R}_2^*)\}. \end{aligned}$$

Specifically, the k -degree meta rule set $\mathcal{P}_{\overline{G}}^{(k)}$ contains all the rules from $(k-1)$ -degree meta rule set $\mathcal{P}_{\overline{G}}^{(k-1)}$ as well as other newly introduced rules $\mathcal{P}^{(k)}$. Each rule contains one non-terminal node \mathcal{R}^* on the *LHS* and two non-terminal nodes \mathcal{R}_1^* and \mathcal{R}_2^* on the *RHS*. The *LHS* of $p_0^{(k)}$ has a degree of $k-1$ as it contains $k-1$ anchor nodes V_{anc} . The *RHS* of $p_0^{(k)}$ attaches

one non-terminal node \mathcal{R}_2^* to the other non-terminal node \mathcal{R}_1^* which connects all the anchor nodes. Therefore, the *RHS* has a degree of k which is larger than the degree of the *LHS* by one. In all the other rules, the *LHS* has k anchor nodes and is of degree k while the *RHS* keeps a maximal degree of k . The k anchor nodes on the *RHS* are combinatorially distributed to \mathcal{R}_1^* and \mathcal{R}_2^* , i.e. if \mathcal{R}_1^* is attached to i anchor nodes, \mathcal{R}_2^* is attached to the rest of $k-i$ anchor nodes. Since we consider \mathcal{R}_1^* and \mathcal{R}_2^* to be the same during grammar production, the range of i is $\{1, \dots, \lfloor \frac{k}{2} \rfloor\}$. In total, there are $1 + \lfloor \frac{k}{2} \rfloor$ rules in $\mathcal{P}^{(k)}$.

Proof of Edit Completeness. According to the definition in (Zhang, 1996; Paaßen, 2018), tree edit distance $TED(T, T')$ between two trees T and T' is the minimum number of operations required to transform one tree into the other. There are three types of available operations: inserting, deleting, and relabeling. Consider two arbitrary trees T and T' that satisfy the definition of edit completeness (being derived from the meta grammar, $TED(T, T') = 1$, and $|T| < |T'|$). Since all the trees generated from the meta grammar have homogeneous tree node labels (i.e. \mathcal{R}^*), there is no relabeling operation in TED calculation. Furthermore, the tree size relation between T and T' admits only one type of tree edit operation: a one-step inserting. Figure 8b illustrates the one-step inserting operation. It operates on a certain node v in T whose subtree is $T_{sub}(v)$. A new node v' , which takes a subset of $T_{sub}(v)$ as the children, is added as a new child of node v . The rest of $T_{sub}(v)$ stays as the children of node v . Suppose the degree of node v in T is m and the degree related to subtree $T_{sub}(v)$ is n . Without loss of generality, we suppose $0 \leq n \leq m/2$. We can treat v in T as the \mathcal{R}^* on the *LHS*, v' and v in T' as \mathcal{R}_1^* and \mathcal{R}_2^* on the *RHS*, respectively. Then the one-step inserting operation can be achieved using the rule $p_n^{(m+1)} \in \mathcal{P}^{(m+1)}$ from $(m+1)$ -degree meta rule set $\mathcal{P}_{\overline{G}}^{(m+1)}$ as constructed above. Since $\Delta(T') \geq m+1$, the meta grammar $\mathcal{P}_{\overline{G}}^{(k)}$ that generates T' has $k \geq m+1$. So $p_n^{(m+1)} \in \mathcal{P}^{(m+1)} \subseteq \mathcal{P}_{\overline{G}}^{(k)}$; thus finish the proof.

Proof of Degree k . For an arbitrary tree T with maximal degree k , we prove that meta grammar $\mathcal{P}_{\overline{G}}^{(k)}$ can generate T by finding a reverse path of production rule from the tree T to the initial node \mathcal{X} . We first find all the nodes $\{v_i\}_{i=1}^n$ in T that have a degree of k . For each node, by edit completeness, we can use the inverse operation of inserting and make the node to have a degree of $k-1$. It can be shown from the proof for edit completeness that each inverse operation of inserting corresponds to an inverse deployment of a production rule. By using n reverse production steps on all the k -degree nodes, we can obtain a T' with $\Delta(T') = k-1$. We can continue the process until the tree is 0-degree, i.e. the \mathcal{X} ; thus finish the proof.

Proof of Minimality. From edit completeness, we know each production rule in the meta grammar corresponds to at least one case of one-step inserting operation. As constrained by degree k , arbitrary trees that are of degree smaller than k need to be covered by the meta grammar. We thus need to handle all the possible cases of one-step inserting operation when finding the reverse production path from the tree to \mathcal{X} . Note that one-step inserting operation involves the division of a subtree, which has combinatorial ways. Then every combinatorial distribution of anchor nodes corresponds to one case of subtree division. This proves each production rule in the meta rule set is necessary. Since there are no duplicate rules, it can be concluded that the meta grammar constructed above is minimal.

4-degree Meta Rules. Figure 8a shows the 4-degree meta grammar we use in our approach. Since the edit completeness is the foundation to prove all the other attributes of the meta grammar, we use a tree edit distance toolkit³ and verify its correctness for the 4-degree meta grammar.

J. Molecular Rules Construction

To obtain the input for the molecular rule construction, we first convert the molecule into a molecular hypergraph $H = (V, E_H)$. A node $v \in V$ represents an atom of the molecule. A hyperedge $e \in E_H$ corresponds to either a bond that joins

³<https://pythonhosted.org/zss/>

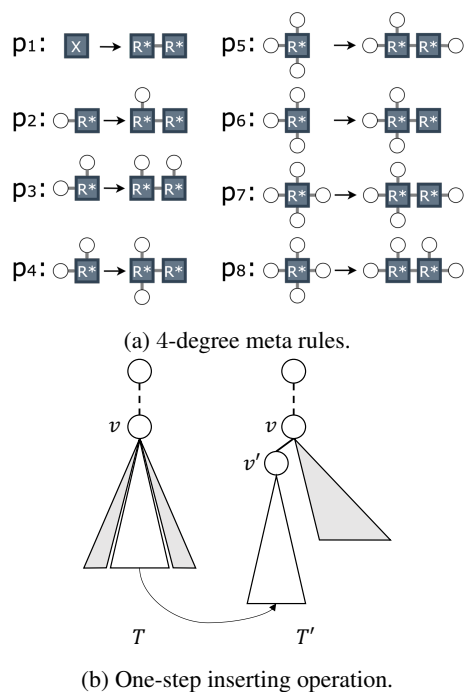


Figure 8: Illustration of 4-degree meta grammar and one-step inserting operation.

Algorithm 1 Molecular Rules Construction.

Input : molecular hypergraph $H = (V, E_H)$, probability function $\phi(\cdot; \theta)$

Output : molecular rule set $\mathcal{P}_{mol} = \{p_i\}_{i=1}^M$, junction tree $T = (V_T, E_T)$ where each $v_T \in V_T$ is a subset of V

- 1 Initialize $\mathcal{P}_{mol} = \emptyset, V_T = \emptyset, E_T = \emptyset$;
 Create a set for unvisited nodes $V_{unv} \leftarrow V$, a set for unvisited hyperedges $E_{unv} \leftarrow E_H$;
- 2 **while** $E_{unv} \neq \emptyset$ **do**
- 3 $\mathbf{e} = [e_1, \dots, e_K], e_i \in E_{unv}, i = 1, \dots, K = |E_{unv}|$;
 $\mathbf{X} \sim \text{Bernoulli}(\phi(\mathbf{e}; \theta))$;
 /* Construct a rule for each connected component */
- 4 **for each** $H_{sub} = (V_{sub}, E_{sub})$ **in** ConnectedComponents(\mathbf{X}, H) **do**
 /* Separate the connected component */
- 5 $\hat{V}_{sub} \leftarrow V_{sub}, \hat{E}_{sub} \leftarrow E_{sub}, V_{anc} \leftarrow \emptyset, E_{anc} \leftarrow \emptyset$;
 for each $v \in V_{sub}$ **do**
- 6 $V_n = \text{GetNeighbors}(v, H)$;
 if $V_n \cap V_{sub} \neq \emptyset$ **then**
- 7 **if** $v \notin V_{unv}$ **or** InRing(v, H) **then**
- 8 $\hat{V}_{sub} \leftarrow \hat{V}_{sub} \setminus \{v\}$;
 $\hat{E}_{sub} \leftarrow \hat{E}_{sub} \setminus \{(s, v) | (s, v) \in E_{sub}\}$;
 $V_{anc} \leftarrow V_{anc} \cup \{v_{anc}^{(i)}, i = |V_{anc}| + 1\}$;
 $E_{anc} \leftarrow E_{anc} \cup \{(s, v_{anc}^{(i)}) | (s, v) \in E_{sub}, i = |V_{anc}| + 1\}$;
- 9 **else**
- 10 $V_{anc} \leftarrow V_{anc} \cup \{v_{anc}^{(i)}, i = |V_{anc}| + k\}_{k=1}^{|V_n \cap V_{sub}|}$;
 $E_{anc} \leftarrow E_{anc} \cup \{(s, v_{anc}^{(i)}) | (s, v) \in E_{sub}, i = |V_{anc}| + k\}_{k=1}^{|V_n \cap V_{sub}|}$;
- 11 /* Construct a production rule */
- 12 $p = \text{RuleConstruction}(\hat{V}_{sub}, \hat{E}_{sub}, V_{anc}, E_{anc})$;
 $\mathcal{P}_{mol} \leftarrow \mathcal{P}_{mol} \cup \{p\}$;
 /* Construct junction tree */
- 13 $v_T = (V_{sub})$;
 for each $\{v_t \in V_T | v_t \cap v_T \neq \emptyset\}$ **do**
- 14 $E_T \leftarrow E_T \cup \{(v_t, v_T)\}$;
- 15 $V_T \leftarrow V_T \cup \{v_T\}$;
 /* update visited status of nodes and hyperedges */
- 16 $V_{unv} \leftarrow V_{unv} \cup V_{sub}, E_{unv} \leftarrow E_{unv} \cup E_{sub}$;
- 17 **return** $\mathcal{P}_{mol}, T = (V_T, E_T)$;

only two nodes or a ring (including aromatic ones) that joins all nodes in the ring. An illustration is provided in Figure 2 of (Guo et al., 2022). The probability function $\phi(\cdot; \theta)$ is defined on each hyperedge: $\phi(e; \theta) = \sigma(-\mathcal{F}_\theta(f(e)))$, where $\sigma(\cdot)$ is the sigmoid function, $\mathcal{F}_\theta(\cdot)$ is a two-layer fully connected network whose final output dimension is 1, and $f(\cdot)$ is a feature extractor using a pre-trained GNN (Hu et al., 2020).

Algorithm 1 illustrates the detailed process of constructing molecular rules for a single molecule. Note that in our approach, the molecular rule construction is performed simultaneously for all input molecules. For each molecule, we first perform an i.i.d. sampling on all the hyperedges following a Bernoulli distribution which takes the value 1 with probability indicated by $\phi(e; \theta)$. We then obtain a binary vector \mathbf{X} that indicates whether each hyperedge is sampled (line 4-5). Next, all connected components are extracted with respect to the sampled hyperedges (line 6). A production rule is constructed for each connected component $H_{sub} = (V_{sub}, E_{sub})$ (line 7-19). Specifically, a production rule needs two components: a hypergraph $\hat{H}_{sub} = (\hat{V}_{sub}, \hat{E}_{sub})$ for *RHS* and anchor nodes V_{anc} indicating the correspondence between the *LHS* and the *RHS*. \hat{H}_{sub} is obtained by removing visited nodes and in-ring nodes from H_{sub} . V_{anc} contains nodes from V that are connected to nodes from H_{sub} but do not appear in H_{sub} themselves. E_{anc} are the edges connecting anchor nodes to \hat{H}_{sub} following the

same connectivity in the original graph H . Following (Guo et al., 2022), the function `RuleConstruction` ($\hat{V}_{sub}, \hat{E}_{sub}, V_{anc}, E_{anc}$) returns a production rule $p : LHS \rightarrow RHS$ constructed as

$$\begin{aligned} LHS &:= H(V_L, E_L), V_L = \{\mathcal{R}^*\} \cup V_{anc}, E_L = \{(\mathcal{R}^*, v) | v \in V_{anc}\}, \\ RHS &:= H(V_R, E_R), V_R = \hat{V}_{sub} \cup V_{anc}, E_R = \hat{E}_{sub} \cup E_{anc}. \end{aligned} \quad (3)$$

In contrast to (Guo et al., 2022), our molecular rule construction does not replace the connected component with the non-terminal node \mathcal{R}^* at every iteration since according to our definition, the molecular rule does not contain any non-terminal nodes on the RHS .

The junction tree of the molecule is constructed along with the construction of molecular rules (line 21-24). We simply treat each connected component V_{sub} sampled at each iteration as a node $v_T \in V_T$ in the junction tree $T = (V_T, E_T)$. Two nodes are connected by an edge $e_T \in E_T$ if their corresponding connected components share hypergraph nodes from H . Since each hyperedge in H is only visited once, the constructed junction tree then contains all the hyperedges and nodes of H without redundancy (Kajino, 2019).

K. Grammar-induced Geometry Construction

Algorithm 2 Grammar-induced Geometry Construction.

Input : meta production rules $\mathcal{P}_{\overline{\mathcal{G}}} = \{p_i\}_{i=1}^N$, maximum BFS depth D , a set of molecular hypergraphs $\{H_i\}_{i=1}^M$ and their corresponding junction trees $\mathcal{T} = \{T_i\}_{i=1}^M$

Output : geometry in the form of a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where each $v = H_v = (V_v, E_v) \in \mathcal{V}$ represents a meta tree or a molecular hypergraph and \mathcal{E} is the edge set of the geometry

```

/* Add root of the geometry */
16 Initialize  $v_{root} = H_{root} = (\mathcal{X}, \emptyset)$ ,  $\mathcal{V} = \{v_{root}\}$ ,  $\mathcal{E} = \emptyset$ ;
    Create a queue data structure  $Q$ ;
     $Q.enqueue(v_{root})$ ;
    /* Breadth-first search for meta geometry, pre-computed offline */
17 while  $Q$  is not empty and  $|\mathcal{V}| \leq D$  do
18      $v = Q.dequeue()$ ;
        /* Expand the meta geometry */
19     for each  $p_i \in \mathcal{P}_{\overline{\mathcal{G}}}$  do
20         if  $p_i$  is applicable to  $H_v$  then
21              $H_v \xrightarrow{p_i} H_{new}$ ;
                 $v_{new} = H_{new}$ ;
                if  $v_{new} \notin \mathcal{V}$  then
22                  $\mathcal{V} \leftarrow \mathcal{V} \cup \{v_{new}\}$ ,  $\mathcal{E} \leftarrow \mathcal{E} \cup \{(v, v_{new})\}$ ;
                     $Q.enqueue(v_{new})$ ;
23             else
24                  $\hat{v} = \text{GetIsomorphicGraph}(H_{new}, \mathcal{V})$ ,  $\mathcal{E} \leftarrow \mathcal{E} \cup \{(v, \hat{v})\}$ ;
/* Connect molecular leaves during run-time */
25 for each  $T_i \in \mathcal{T}$  do
26     for each  $v \in \mathcal{V}$  do
27         if  $\text{IsIsomorphic}(H_v, T_i)$  then
28              $v_i = H_i$ ;
                 $\mathcal{V} \leftarrow \mathcal{V} \cup \{v_i\}$ ,  $\mathcal{E} \leftarrow \mathcal{E} \cup \{(v, v_i)\}$ ;
29 return  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ ;

```

Algorithm 2 illustrates the detailed algorithm to construct the grammar-induced geometry. It contains two parts: the construction of the meta geometry (line 4-14) and the construction of the molecular leaves (line 15-19). The meta geometry construction follows the standard breadth-first search (BFS) starting from the root $H_{root} = (\mathcal{X}, \emptyset)$ (line 1). Every time

when we visit a node $v \in \mathcal{V}$ in \mathcal{G} , we find all the rules that are applicable to H_v from the meta rule set (line 6-7). Each applicable rule is applied to H_v in order to create a new meta tree H_{new} (line 8). Depending on whether there is a node in \mathcal{V} that represents a isomorphic tree to H_{new} in the current geometry, we either create a new node v_{new} (line 9) or find the existing matched node \hat{v} (line 14). We then add an edge between v and v_{new} , or between v and \hat{v} (line 11 and 14). The function `GetIsomorphicGraph(H_{new}, \mathcal{V})` enumerates every node in \mathcal{V} , checks if the tree represented by the node is isomorphic to H_{new} , and returns the matched node. We use the algorithm from (Cordella et al., 2001) which is implemented in `networkx`⁴ for graph isomorphism test. To increase the speed of `GetIsomorphicGraph(H_{new}, \mathcal{V})`, we use Weisfeiler Lehman graph hash (Shervashidze et al., 2011) and only perform isomorphism test for graph pairs that share the same hashing code.

The molecular leaves are added during run-time (line 15-19). For each input molecular hypergraph, we check if there is a meta tree node that is isomorphic to its junction tree. If so, we add an edge connecting the matched meta tree node to the molecular hypergraph. The function `IsIsomorphic(H_v, T_i)` is implemented using the same package for `GetIsomorphicGraph(H_{new}, \mathcal{V})` and returns true if H_v and T_i are isomorphic to each other.

L. Optimization

Figure 9 shows a graphical model of dependency in our optimization problem. The two sets of parameters θ and (φ, ψ, α) are independent and thus can be iteratively optimized using block coordinate descent. Since the objective in Equation 2 is differentiable with respect to (φ, ψ, α) , we can use gradient descent for the optimization iterations related to the graph neural diffusion. The geometry learning, however, is non-differentiable due to the fact that we perform sampling to construct molecular rules and construct the geometry. Hence, we rewrite the objective for optimizing θ in an expectation form and apply REINFORCE (Williams, 1992) to obtain a stochastic gradient, as done in (Guo et al., 2022):

$$\begin{aligned} \min_{\theta} l(\mathbf{u}_T, \hat{\mathbf{u}}) &= \min_{\theta} \mathbb{E}_{\mathbf{X}} [l(\mathbf{u}_T, \hat{\mathbf{u}})], \\ \nabla_{\theta} \mathbb{E}_{\mathbf{X}} [l(\mathbf{u}_T, \hat{\mathbf{u}})] &= \int_{\mathbf{X}} l(\mathbf{u}_T, \hat{\mathbf{u}}) \nabla_{\theta} p(\mathbf{X}) \\ &= \mathbb{E}_{\mathbf{X}} [l(\mathbf{u}_T, \hat{\mathbf{u}}) \nabla_{\theta} \log(p(\mathbf{X}))] \approx \frac{1}{N} \sum_{n=1}^N l(\mathbf{u}_T^{(n)}, \hat{\mathbf{u}}) \nabla_{\theta} \log(p(\mathbf{X}^{(n)})), \end{aligned}$$

where \mathbf{X} is a concatenation of binary vectors indicating how hyperedges are sampled in molecular rule construction in Appendix J.

M. Examples of Retro-Synthesis Paths

Figure 10 shows the retro-synthesis paths of three molecules generated using our pipeline. Since Retro* score is used as one of the grammar metrics, our approach is capable of retro-synthesis planning for all the generated molecules, providing a complete pipeline for novel molecule discovery.




Figure 9: A graphical model of dependency.

⁴<https://networkx.org/documentation/stable/reference/algorithms/isomorphism.html>

