

---

# Molecular Hypergraph Grammar with Its Application to Molecular Optimization

---

Hiroshi Kajino<sup>1</sup>

## Abstract

Molecular optimization aims to discover novel molecules with desirable properties, and its two fundamental challenges are: (i) it is not trivial to generate valid molecules in a controllable way due to hard chemical constraints such as the valency conditions, and (ii) it is often costly to evaluate a property of a novel molecule, and therefore, the number of property evaluations is limited. These challenges are to some extent alleviated by a combination of a variational autoencoder (VAE) and Bayesian optimization (BO), where VAE converts a molecule into/from its latent continuous vector, and BO optimizes a latent continuous vector (and its corresponding molecule) within a limited number of property evaluations. While the most recent work, for the first time, achieved 100% validity, its architecture is rather complex due to auxiliary neural networks other than VAE, making it difficult to train. This paper presents a *molecular hypergraph grammar variational autoencoder* (MHG-VAE), which uses a single VAE to achieve 100% validity. Our idea is to develop a graph grammar encoding the hard chemical constraints, called *molecular hypergraph grammar* (MHG), which guides VAE to always generate valid molecules. We also present an algorithm to construct MHG from a set of molecules.

## 1. Introduction

Molecular optimization aims to discover a novel molecule that possesses prescribed properties given by a user. For example, Gómez-Bombarelli et al. (2016) aim to maximize the efficiency of an organic light-emitting diode. Letting  $\mathcal{M}$  be a set of valid molecules, the molecular optimization

<sup>1</sup>MIT-IBM Watson AI Lab; IBM Research, Tokyo, Japan. Correspondence to: Hiroshi Kajino <kajino@jp.ibm.com>.

problem is to obtain  $m^* \in \mathcal{M}$  such that,

$$m^* = \arg \max_{m \in \mathcal{M}} f(m), \quad (1)$$

where  $f: \mathcal{M} \rightarrow \mathbb{R}$  is an unknown function that outputs a chemical property of the input molecule to be maximized.

There are two obstacles to solve Eq. (1). First, the set of feasible solutions  $\mathcal{M}$  is discrete, and it is difficult to generate a candidate  $m$  from  $\mathcal{M}$  in a controllable way. Second,  $f$  is unknown and costly to evaluate, and only a small sample  $\{(m_n, y_n)\}_{n=1}^N \subset \mathcal{M} \times \mathbb{R}$  is available, where  $y_n = f(m_n)$ . In fact, the function evaluation often requires wet-lab experiments or day-long simulation based on quantum mechanics.

A recent innovation (Gómez-Bombarelli et al., 2018) facilitates the optimization by leveraging VAE (Kingma & Welling, 2014) and BO (Moćkus, 1975). The first challenge is addressed by casting the discrete optimization problem into continuous with the help of VAE. In specific, they first train a VAE, a pair of Enc:  $\mathcal{M} \rightarrow \mathbb{R}^D$  and Dec:  $\mathbb{R}^D \rightarrow \mathcal{M}$  such that  $\text{Dec}(\text{Enc}(m)) \approx m$  holds for any  $m \in \mathcal{M}$ , and they obtain  $m^*$  as follows:

$$m^* = \text{Dec} \left( \arg \max_{z \in \mathbb{R}^D} f(\text{Dec}(z)) \right), \quad (2)$$

The second challenge is addressed by using BO, which iteratively optimizes a black-box function in a limited number of function evaluations. In specific, at each iteration, the inner optimization problem in Eq. (2) is solved by BO with training sample  $\{(\text{Enc}(m_1), y_1), \dots, (\text{Enc}(m_N), y_N)\} \subset \mathbb{R}^D \times \mathbb{R}$ , and the resultant latent vector  $x_{N+1} \in \mathbb{R}^D$  is added to the sample with property evaluation  $f(\text{Dec}(x_{N+1}))$ . This procedure is iterated for a fixed number of times, and BO outputs the best molecule found so far.

While they elegantly address the two obstacles, the decoding sometimes fails, and no molecule is obtained, which we call the *decoding error issue*. This is mainly due to the use of SMILES (Weininger, 1988) to represent a molecule. Let  $\Sigma$  be a set of symbols used in SMILES, and  $\text{Enc}_\Sigma: \mathcal{M} \rightarrow \Sigma^*$  be a SMILES encoder. For example,  $\Sigma$  includes atomic symbols, e.g., C, H  $\in \Sigma$ ; given a phenol as input, the encoder outputs c1cc(O)cccc1, where the digits represent

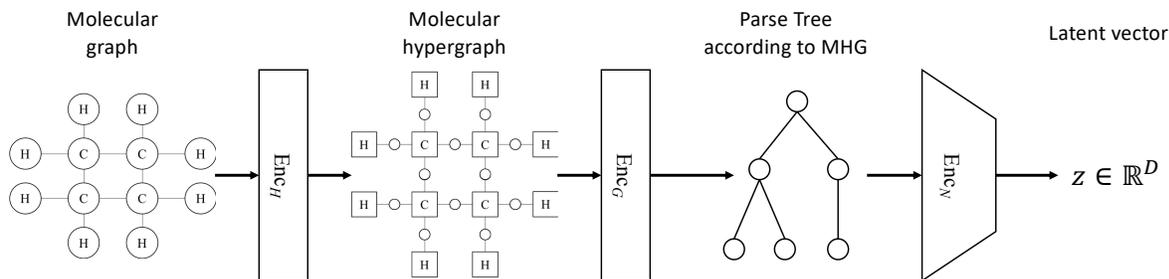


Figure 1. Illustration of our encoder. For a molecular hypergraph, squares represent hyperedges, circles represent nodes, and a circle-square line indicates that the node is a member of the hyperedge. The decoder is defined by inverting the encoder.

the start and end points of the benzene ring, the parentheses represent branching, and hydrogen atoms are omitted. Letting  $\text{Enc}_S[\mathcal{M}] := \{\text{Enc}_S(m) \mid m \in \mathcal{M}\} \subsetneq \Sigma^*$  be the set of all *valid* SMILES strings, a SMILES decoder  $\text{Dec}_S: \text{Enc}_S[\mathcal{M}] \rightarrow \mathcal{M}$  can be defined. Notice that the domain of  $\text{Dec}_S$  is not  $\Sigma^*$  but  $\text{Enc}_S[\mathcal{M}]$ , the set of strings that follow SMILES’ grammar, because any string that violates the grammar cannot be decoded into any molecule.

In their implementation, the encoder is composed as  $\text{Enc} = \text{Enc}_N \circ \text{Enc}_S$ , where  $\text{Enc}_N: \Sigma^* \rightarrow \mathbb{R}^D$  is a neural network encoder, and the decoder is composed as  $\text{Dec} = \text{Dec}_S \circ \text{Dec}_N$ , where  $\text{Dec}_N: \mathbb{R}^D \rightarrow \Sigma^*$  is a neural network decoder, generating symbols one by one. The decoding error issue occurs when the output of  $\text{Dec}_N$  does not belong to  $\text{Enc}_S[\mathcal{M}]$ , the domain of  $\text{Dec}_S$ . For the phenol example, if  $\text{Dec}_N$  fails to output the end digit, c1c(O)cccc, it cannot be converted into a molecule because the ring cannot be closed; if  $\text{Dec}_N$  generates more than one pair of parentheses, c1c(O)(O)cccc1, it violates the valence condition of carbon. Since SMILES’ grammar is a context-sensitive grammar, it is not straightforward to develop a neural network that always generates a string belonging to  $\text{Enc}_S[\mathcal{M}]$ .

Recently, several studies have been conducted towards addressing the decoding error issue (Kusner et al., 2017; Dai et al., 2018; Jin et al., 2018). Among them, Jin et al. (2018) for the first time report 100% validity, addressing the decoding error issue. Their idea is to represent a molecular graph as fragments (such as rings and atom branches) connected in a tree structure. Such a tree representation is preferable because it is easier to generate a tree than a general graph with degree constraints. By forcing the decoder to generate only valid combination of fragments, the decoder can always generate a valid molecule.

While their tree representation successfully addresses the decoding error issue, it models only part of molecular properties, and the rest is left to neural networks. For example, their representation only specifies fragment-level connections, and does not specify which atoms in the fragments to be connected. In addition, since it does not specify atom-

level connections, the stereochemistry information disappears. They instead enumerate all possible configurations and pick one by training several auxiliary neural networks.

Given this literature, we are interested in addressing the decoding error issue *without* the auxiliary neural networks. Such a simple architecture will facilitate model training, and it will be easier to achieve high performance with less effort. Our idea is to develop a (i) context-free graph grammar of a molecular graph that (ii) never generates an invalid molecule, and that (iii) can represent the atom-level connection and stereochemistry information. With such a grammar, we can use a parse tree as an intermediate representation of a molecule, which is easy for the decoder to generate (as guaranteed by the context-freeness (i)). Combined with (ii) and (iii), our decoder can always output a valid molecule with stereochemistry information, using a single VAE only.

Our technical highlight is a *molecular hypergraph grammar* (MHG) along with an algorithm to infer MHG from a set of molecules. MHG is designed to satisfy all of the requirements mentioned above. We also develop a molecular hypergraph grammar variational autoencoder (MHG-VAE), which combines MHG and VAE to obtain an autoencoder for a molecule. Figure 1 illustrates our encoder, where the first two encoders,  $\text{Enc}_H$  and  $\text{Enc}_G$ , and their corresponding decoders  $\text{Dec}_H$  and  $\text{Dec}_G$ , are our main contributions, while we use a standard seq2seq VAE for  $(\text{Enc}_N, \text{Dec}_N)$ .

In details, we develop MHG by tailoring a hyperedge replacement grammar (HRG) (Drewes et al., 1997) for a *molecular hypergraph* (thus, MHG is a special case of HRG). A molecular hypergraph models an atom by a hyperedge and a bond by a node. HRG is a context-free graph grammar generating a hypergraph by replacing a non-terminal hyperedge with another hypergraph; it achieves atom-level connections when combined with a molecular hypergraph, and stereochemistry can also be encoded into the grammar. It also preserves the number of nodes belonging to each hyperedge, which coincides with the valency of an atom in our case. Therefore, these two ideas allow us to always generate valid molecules using a single VAE.

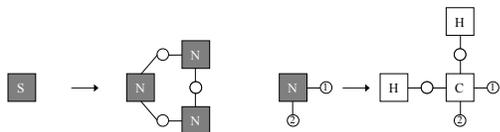


Figure 2. Part of production rules extracted from the hypergraph in Fig. 1, illustrated as  $A \rightarrow R$ . Filled squares represent non-terminals, and unfilled ones terminals. The numbered nodes are external nodes, where the numbers indicate the correspondence between the nodes in  $A$  and  $R$  when  $A$  is replaced with  $R$ .

Our MHG inference algorithm extends the existing HRG inference algorithm (Aguinaga et al., 2016) so that the resultant HRG always generates a molecular hypergraph. The existing one infers HRG by extracting a set of production rules from a tree decomposition of each hypergraph, which is equivalent to a parse tree. Our finding is that, while the inferred HRG preserves the valence condition, it sometimes generates a hypergraph that cannot be decoded into a molecular graph; the generated hypergraph may contain a node that is shared by more than two hyperedges. To address this issue, we develop an *irredundant* tree decomposition, with which HRG is guaranteed to generate a valid molecular hypergraph, *i.e.*, the inferred HRG is always MHG.

## 2. Preliminaries

A hypergraph is a pair  $H = (V_H, E_H)$ , where  $V_H$  is a set of nodes, and  $E_H$  is a set of non-empty subsets of  $V_H$ , called hyperedges. A hypergraph is called  $k$ -regular if every node has degree  $k$ .

A tree decomposition of a hypergraph (Def. 1) discovers a tree-like structure of the hypergraph. Figure 3 illustrates a tree decomposition of the hypergraph shown in Fig. 1.

**Definition 1.** A tree decomposition of hypergraph  $H = (V_H, E_H)$  is tree  $T = (V_T, E_T)$  with two labeling functions  $\ell_T^{(V)}: V_T \rightarrow 2^{V_H}$  and  $\ell_T^{(E)}: V_T \rightarrow 2^{E_H}$  such that:

1. For each  $v_H \in V_H$ , there exists at least one node  $v_T \in V_T$  such that  $v_H \in \ell_T^{(V)}(v_T)$ .
2. For each  $e_H \in E_H$ , there exists exactly one node  $v_T \in V_T$  such that  $e_H \subseteq \ell_T^{(V)}(v_T)$  and  $e_H \in \ell_T^{(E)}(v_T)$ .
3. For each  $v_H \in V_H$ , a set of nodes  $\{v_T \in V_T \mid v_H \in \ell_T^{(V)}(v_T)\}$  is connected in  $T$ .

Let us denote the hypergraph on node  $v_T \in V_T$  by  $H(v_T) := (\ell_T^{(V)}(v_T), \ell_T^{(E)}(v_T))$ .

A hyperedge replacement grammar (HRG) (Drewes et al., 1997) is a context-free grammar generating hypergraphs with labeled nodes and hyperedges (Def. 2, Fig. 2). It starts

from the starting symbol  $S$  and repeatedly replaces a non-terminal symbol  $A$  in the hypergraph with a hypergraph  $R$ , which may have both terminal and non-terminal symbols.

**Definition 2.** A hyperedge replacement grammar is a tuple  $\text{HRG} = (N, T, S, P)$ , where,

1.  $N$  is a set of non-terminal hyperedge labels.
2.  $T$  is a set of terminal hyperedge labels.
3.  $S \in N$  is the starting non-terminal hyperedge.
4.  $P$  is a set of production rules where,
  - $p = (A, R) \in P$  is a production rule,
  - $A \in N$  is a non-terminal symbol, and
  - $R$  is a hypergraph with hyperedge labels  $T \cup N$  and has  $|A|$  external nodes.

We define a *parse tree* according to HRG as follows. Each node of the parse tree is labeled by a production rule. The production rules of the leaves of the parse tree must not contain non-terminals in their  $R$ s. If the production rule  $p$  is a starting rule, the node has  $N_p$  ordered children, where  $N_p$  denotes the number of non-terminals in  $R$ , and the edges are ordered by the orders of the non-terminals. Otherwise, the node has one parent and  $N_p$  ordered children, where the corresponding non-terminal in the parental production rule must coincide with  $A$  of the production rule.

Given a parse tree, we can construct a hypergraph by sequentially applying the production rules. Such sequential applications of production rules are equivalent to the parse tree, and we call it a *parse sequence*.

## 3. Molecular Graph and Hypergraph

This section introduces our definitions of a molecular graph and a molecular hypergraph. We also present a pair of encoder and decoder between them,  $(\text{Enc}_H, \text{Dec}_H)$ .

### 3.1. Molecular Graph

A molecular graph (Def. 3) represents the structural formula of a molecule using a graph, where atoms are modeled as labeled nodes and bonds as labeled edges. Typically, the node label is defined by the atom’s symbol (*e.g.*, H, C) and its formal charge, and the edge label by the bond type (*e.g.*, single, double). The graph must satisfy the valency condition; the degree of each atom is specified by its label (*e.g.*, the degree of C must equal four). Let  $\mathcal{G}(L_G^{(V)}, L_G^{(E)}, d^{(V)})$  be the set of all possible molecular graphs, given the sets of node and edge labels and the degree constraint function.

**Definition 3.** Let  $L_G^{(V)}$  and  $L_G^{(E)}$  be sets of node and edge labels. Let  $d^{(V)}: L_G^{(V)} \rightarrow \mathbb{N}$  be a degree constraint function.

Let  $G = (V_G, E_G, \ell_G^{(V)}, \ell_G^{(E)})$  be a node and edge-labeled graph, where  $V_G$  is a set of nodes,  $E_G$  is a set of undirected edges,  $\ell_G^{(V)}: V_G \rightarrow L_G^{(V)}$  is a node-labeling function, and  $\ell_G^{(E)}: E_G \rightarrow L_G^{(E)}$  is an edge-labeling function. A molecular graph  $G$  is a node and edge-labeled graph that satisfies  $d(v) = d^{(V)}(\ell_G^{(V)}(v))$  for all  $v \in V_G$ , where  $d(v)$  indicates the degree of node  $v$ .

There are two types of important properties that influence the chemical properties of a molecule. The first one is the aromaticity of a ring (e.g., benzene derivatives). The bonds in an aromatic ring are different from a single or double bond, and are known to be more stable. We do not explicitly encode any information related to the aromaticity, and instead, employ the Kekulé structure, where an aromatic ring is represented by alternating single and double bonds. This does not lose generality because we can infer the aromaticity from the Kekulé representation. The second one is the stereochemistry, which specifies 2D or 3D configuration of atoms. We deal with the configuration at a double bond and tetrahedral carbon. The double bond configuration is encoded by an E-Z configuration label assigned on the edge label. Given the label and the whole structure of the molecule, the Cahn–Ingold–Prelog priority rules can specify the double bond direction. For the tetrahedral chirality information, we assign a chirality tag in the node label, following the implementation of RDKit.

In summary, we employ the graph representation (Def. 3), where the node label contains the atom symbol, formal charge, and the tetrahedral chirality tag, and the edge label contains the bond type and the E-Z configuration.

### 3.2. Molecular Hypergraph

As an intermediate representation, we use a molecular hypergraph (Def. 4), where an atom is modeled by a hyperedge and a bond between two atoms by a node shared by the corresponding two hyperedges. Let  $\mathcal{H}(L_H^{(V)}, L_H^{(E)}, c^{(E)})$  be the set of all molecular hypergraphs, given the sets of node and hyperedge labels and the cardinality constraint function.

**Definition 4.** Let  $L_H^{(E)}$  and  $L_H^{(V)}$  be sets of hyperedge and node labels. Let  $c^{(E)}: L_H^{(E)} \rightarrow \mathbb{N}$  be a cardinality constraint function. Let  $H = (V_H, E_H, \ell_H^{(E)}, \ell_H^{(V)})$  be a node and hyperedge-labeled hypergraph, where  $V_H$  is a set of nodes,  $E_H$  is a set of hyperedges,  $\ell_H^{(V)}: V_H \rightarrow L_H^{(V)}$  is a node-labeling function, and  $\ell_H^{(E)}: E_H \rightarrow L_H^{(E)}$  is a hyperedge-labeling function. A molecular hypergraph  $H$  is a node and hyperedge-labeled hypergraph that satisfies the followings:

1. (Regularity)  $H$  is 2-regular.
2. (Cardinality) for each  $e \in E_H$ ,  $|e| = c^{(E)}(\ell_H^{(E)}(e))$  holds, where  $|e|$  is the cardinality of hyperedge  $e$ .

Note that the regularity condition in Def. 4 assures that any molecular hypergraph can be decoded into a graph.

### 3.3. Encoder and Decoder

Finally, we present the encoder and decoder between a molecular graph and a molecular hypergraph,  $(\text{Enc}_H, \text{Dec}_H)$ . They can be derived easily by swapping nodes-hyperedges and edges-nodes. The regularity condition in Def. 4 assures the swap to work. This equivalence immediately yields the following:

**Theorem 1.** If  $L_G^{(V)} = L_H^{(E)}$ ,  $L_G^{(E)} = L_H^{(V)}$ , and  $d^{(V)}(l) = c^{(E)}(l)$  for all  $l \in L_G^{(V)}$  hold, then the followings hold:

$$\begin{aligned} \mathcal{H}(L_H^{(V)}, L_H^{(E)}, c^{(E)}) &= \text{Enc}_H[\mathcal{G}(L_G^{(V)}, L_G^{(E)}, d^{(V)})], \\ \mathcal{G}(L_G^{(V)}, L_G^{(E)}, d^{(V)}) &= \text{Dec}_H[\mathcal{H}(L_H^{(V)}, L_H^{(E)}, c^{(E)})], \\ G &= \text{Dec}_H(\text{Enc}_H(G)) \quad (\forall G \in \mathcal{G}(L_G^{(V)}, L_G^{(E)}, d^{(V)})). \end{aligned}$$

## 4. Molecular Hypergraph Grammar

A molecular hypergraph grammar (MHG) is defined as an HRG that always generates molecular hypergraphs. Given MHG, we can define  $(\text{Enc}_G, \text{Dec}_G)$  that leverages MHG to represent a molecular hypergraph as a parse sequence.

Let  $\text{MHG} = (N, T, S, P)$  be a molecular hypergraph grammar, and  $\mathcal{L}_{\text{MHG}}$  be its language, i.e., the set of molecular hypergraphs that can be generated by MHG. The encoder  $\text{Enc}_G: \mathcal{L}_{\text{MHG}} \rightarrow P^*$  maps a molecular hypergraph into the corresponding parse sequence. The decoder maps a parse sequence into a molecular hypergraph by sequentially applying the production rules. It accepts a sequence of production rules obtained from  $\mathcal{L}_{\text{MHG}}$  only, because other sequences cannot generate a molecular hypergraph. Thus, the domain of the decoder is defined as  $\text{Dec}_G: \text{Enc}_G[\mathcal{L}_{\text{MHG}}] \rightarrow \mathcal{L}_{\text{MHG}}$ . Clearly, for any  $H \in \mathcal{L}_{\text{MHG}}$ ,  $H = \text{Dec}_G(\text{Enc}_G(H))$  holds.

## 5. MHG Inference Algorithm

We present an algorithm to infer MHG from a set of molecular hypergraphs. Our algorithm (Sec. 5.2) extends an existing HRG inference algorithm (Sec. 5.1), which extracts a set of production rules from tree decompositions of hypergraphs. We need to tailor a novel inference algorithm for MHG because HRG inferred by applying the existing one to molecular hypergraphs is not MHG. The inferred HRG does not necessarily generate a molecular hypergraph because it sometimes violates the regularity condition in Def. 4.

### 5.1. Existing HRG Inference Algorithm

Aguíñaga et al. (2016) propose an algorithm to infer HRG from a set of hypergraphs. Their key observation is that tree decompositions of hypergraphs yield HRG whose as-

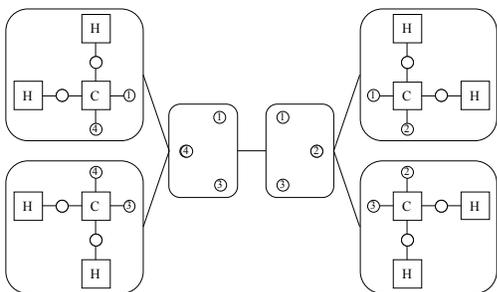


Figure 3. Irredundant tree decomposition of the hypergraph in Fig 1.

sociated language includes the whole input hypergraphs. Assume that we have a tree decomposition  $T$  of hypergraph  $H$ . We arbitrarily choose one node from  $T$  as the root node. For node  $v_T \in V_T$ , let  $\text{pa}(v_T)$  be the parent of  $v_T$  and  $\text{ch}(v_T)$  be a set of children of  $v_T$ . They first notice that connecting each pair  $(v_T, \text{pa}(v_T))$  by their common nodes yields the original hypergraph (e.g., connecting such pairs in Fig. 3 yields the hypergraph in Fig. 1). In other words, a tree decomposition with an arbitrary root node is equivalent to a parse tree. Given this observation, their algorithm extracts a production rule from a triplet  $(\text{pa}(v_T), v_T, \text{ch}(v_T))$  so that the production rule can paste  $H(v_T)$  on  $H(\text{pa}(v_T))$  with non-terminals left for applying the following production rules obtained from the children. Note that the algorithm outputs not only HRG but also parse sequences of input hypergraphs. For more details, see Appendix A.

## 5.2. Our MHG Inference Algorithm

We find that the existing algorithm cannot infer MHG, i.e., the inferred HRG sometimes violates the regularity condition. We develop an *irredundant* tree decomposition so that the violation does not occur, and substitute it for a generic tree decomposition to derive our MHG inference algorithm.

**Irredundant Tree Decomposition.** We introduce a key property of a tree decomposition called *irredundancy*, which is necessary to guarantee the regularity. Intuitively, a tree decomposition is irredundant if each node of the tree does not contain redundant nodes of the original hypergraph. Figures 3 and 4 illustrate both irredundant and redundant tree decompositions. The formal definition appears in Def. 5.

**Definition 5.** Let  $H = (V_H, E_H)$  be a hypergraph, and  $(T, \ell_T^{(V)}, \ell_T^{(E)})$  be its tree decomposition. Let  $V_T(v_H) = \{v_T \in V_T \mid v_H \in \ell_T^{(V)}(v_T)\}$  be a set of nodes in  $T$  that contain  $v_H \in V_H$ . A tree decomposition is irredundant if,

$$\forall v_H \in V_H, \forall v_T \in V_T(v_H) \\ \ell_T^{(E)}(v_T) \neq \emptyset \Leftrightarrow v_T \text{ is a leaf in } T[V_T(v_H)], \quad (3)$$

where  $T[V_T(v_H)]$  is the subgraph induced by  $V_T(v_H)$ .

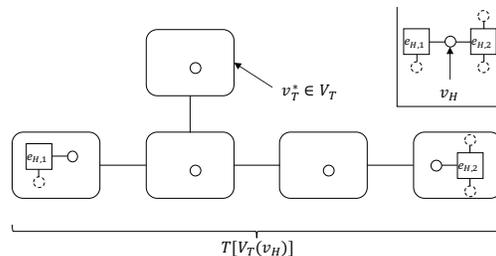


Figure 4. Redundant tree decomposition  $T$ . The upper-right shows a subhypergraph of  $H$ , whose tree decomposition induced by  $V_T(v_H)$  is shown in the lower-left. It is necessary to remove  $v_H$  from  $v_T^*$  to transform it to be irredundant.

We can make any tree decomposition to be irredundant in polynomial time; for each  $v_H$ , if it does not satisfy the condition (3), remove  $v_H$  from each  $\ell_T^{(E)}(v_T)$  ( $v_T \in V_T(v_H)$ ).

**Implementation.** We tailor a tree decomposition algorithm for a molecular hypergraph. It starts from a one-node tree whose node contains the input hypergraph, and updates the tree by applying the following two steps.

The first step is to find a node such that the input hypergraph becomes disjoint when divided at the node, and to divide the hypergraph into two. When dividing a hypergraph at a node, the node is duplicated so that the two hyperedges that the node belonged to still contain the node. This operation is repeatedly applied to the subhypergraphs until there does not exist such a node. As a result, the input hypergraph is divided into (i) hypergraphs containing exactly one hyperedge and (ii) hypergraphs that contain rings. The type-(i) hypergraphs are obtained from branching structures of the input hypergraph, and the type-(ii) ones from rings.

The second step *rips off* the hyperedges from the type-(ii) hypergraphs; we divide a tree node containing a ring into one tree node containing all of the nodes and other tree nodes each of which contains each of the hyperedges and their nodes. This operation is helpful to reduce the number of production rules. The inferred grammar can generate a ring structure by first generating a skeleton of a ring, where all of the hyperedges are non-terminal, and then replacing the non-terminals in an arbitrary way. Without the second step, since the number of possible atom configurations of a ring is enormous, that of production rules also greatly increases.

**Theoretical Result.** Theorem 2 summarizes the properties of our algorithm. It suggests that (i) HRG inferred by our algorithm can generate the whole input hypergraphs and (ii) the inferred HRG always generates a molecular hypergraph, i.e., the HRG is MHG. See Appendix B for its proof.

**Theorem 2.** Let  $\mathcal{H}(L_H^{(V)}, L_H^{(E)}, c^{(E)})$  be a set of all molecular hypergraphs and  $\hat{\mathcal{H}}$  be its finite subset. Let  $\mathcal{L}_{\text{HRG}}(\hat{\mathcal{H}})$

**Algorithm 1** Latent Representation Inference

**In:** Mols and targets,  $\mathcal{G}_0 = \{g_n\}_{n=1}^N$ ,  $\mathcal{Y}_0 = \{y_n\}_{n=1}^N$ .

- 1: Obtain molecular hypergraphs:  $\mathcal{H}_0 \leftarrow \text{Enc}_H(\mathcal{G}_0)$
- 2: Obtain an MHG as well as parse sequences:  
 $\text{Dec}_G, \text{MHG}, \mathcal{S}_0 \leftarrow \text{MHG-Inf}(\mathcal{H}_0)$ .
- 3: Train NNs using  $(\mathcal{S}_0, \mathcal{Y}_0)$  to obtain  $(\text{Enc}_N, \text{Dec}_N)$ ,  $\hat{f}$ .
- 4: Obtain latent vectors,  $\mathcal{Z}_0 \leftarrow \mathbb{E}[\text{Enc}_N[\mathcal{S}_0]]$ .

**return** latent vectors  $\mathcal{Z}_0, \text{Dec}_G, \text{Dec}_N, \hat{f}$ .

be the language generated by HRG inferred by applying our algorithm to  $\hat{\mathcal{H}}$ . Then,  $\mathcal{L}_{\text{HRG}}(\hat{\mathcal{H}})$  satisfies

$$\hat{\mathcal{H}} \subseteq \mathcal{L}_{\text{HRG}}(\hat{\mathcal{H}}) \subseteq \mathcal{H}(L_H^{(V)}, L_H^{(E)}, c^{(E)}).$$

## 6. Application to Molecular Optimization

We have presented  $(\text{Enc}_H, \text{Dec}_H)$ , a pair of encoder and decoder between a molecular graph and a molecular hypergraph (Sec. 3.3), and  $(\text{Enc}_G, \text{Dec}_G)$ , that between a molecular hypergraph and a parse sequence according to MHG (Sec. 4). We have also provided an algorithm to obtain  $(\text{Enc}_G, \text{Dec}_G)$  from a set of molecular hypergraphs (Sec. 5). In this section, we finally present an application of our encoders and decoders to molecular optimization, which aims to search a molecule with desirable properties.

### 6.1. Our Model

Our model consists of encoder and decoder,  $(\text{Enc}, \text{Dec})$ , and optionally, a predictive model from the latent space to a target value. Our encoder and decoder are composed as:

$$\begin{aligned} \text{Enc} &= \text{Enc}_N \circ \text{Enc}_G \circ \text{Enc}_H, \\ \text{Dec} &= \text{Dec}_H \circ \text{Dec}_G \circ \text{Dec}_N, \end{aligned}$$

where  $(\text{Enc}_N, \text{Dec}_N)$  is a seq2seq GVAE (Kusner et al., 2017). Since GVAE can output a parse sequence that follows a context-free grammar,  $\text{Dec}_N$  is guaranteed to output a valid parse sequence that belongs to  $\text{Enc}_G[\mathcal{L}_{\text{MHG}}]$ , the domain of the following decoder  $\text{Dec}_G$ . Our model configuration appears in Appendix C.

### 6.2. Molecular Optimization Algorithm

Global molecular optimization (Eq. (2)) aims to find novel molecules with desirable properties from the entire molecular space. It consists of two steps. Algorithm 1 summarizes the former step to obtain latent representations of the input molecules. Note that this algorithm does not output encoders, because molecular optimization algorithms require Dec only. Algorithm 2 describes the latter step, which optimizes a molecule in the latent space using BO.

**Algorithm 2** Global Molecular Optimization

**In:**  $\mathcal{Z}_0, \mathcal{Y}_0, \text{Dec}, \# \text{iterations } K, \# \text{candidates } M$ .

- 1:  $\mathcal{D}_1 \leftarrow \{(z_n, y_n)\}_{n=1}^N$
- 2: **for**  $k = 1, \dots, K$  **do**
- 3: Fit GP using  $\mathcal{D}_k$ .
- 4: Obtain candidates  $\mathcal{Z}_k = \{z_m \in \mathbb{R}^D\}_{m=1}^M$  from BO.
- 5: Obtain molecular graphs as  $\mathcal{G}_k \leftarrow \text{Dec}[\mathcal{Z}_k]$ .
- 6: Obtain target values as  $\mathcal{Y}_k \leftarrow f[\mathcal{G}_k]$ .
- 7:  $\mathcal{D}_{k+1} \leftarrow \mathcal{D}_k \cup \{(z_{k,m}, y_{k,m}) \in \mathcal{Z}_k \times \mathcal{Y}_k\}_{m=1}^M$ .

**return** novel molecules  $\{g_{k,m}\}_{m=1,k=1}^{M,K}$ .

## 7. Related Work

Molecular optimization has a longstanding history especially in drug discovery, and mostly, combinatorial methods have been used to find novel molecules (Jorgensen, 2009). The paper by Gómez-Bombarelli et al. (2018), for the first time, applies modern machine learning techniques to this problem, and since then, growing number of papers have been tackling this problem.

There are two complementary approaches to molecular optimization. One is the combination of VAE and Bayesian optimization, originally developed by Gómez-Bombarelli et al. (2018), and the other is reinforcement learning, where the construction of a molecule is modeled as a Markov decision process (Guimaraes et al., 2017; You et al., 2018; Zhou et al., 2018). These two approaches cover complementary application areas, and therefore, they are not conflicting with each other. The key difference is the assumption on function evaluation cost. The former assumes that the cost is so high that the number of function evaluations should be kept at a minimum, while the latter assumes that the cost is negligible and a number of trial-and-errors are allowed. Therefore, the former is more favorable when the evaluation requires wet-lab experiments or computationally heavy first-principles calculation, and the latter is more favorable when the evaluation can be carried out by light-weight computer simulation. Since our paper focuses on the former setting, in the following, we will discuss its literature.

As stated in the introduction, the decoding error issue has been one of the critical issues, and therefore, this section focuses on a series of studies alleviating it. There are mainly two approaches to address it. One approach is to devise the decoding network to generate as valid SMILES strings as possible. For example, Kusner et al. (2017) leverage SMILES’ grammar to force the decoder to align the grammar. This approach is limited because they assume a context-free grammar (CFG), while SMILES’ grammar is not totally context-free. Dai et al. (2018) propose to use an attribute grammar, which enhances CFG by introducing attributes and rules. This enhancement allows us to enforce semantic constraints to the decoder. However, they deal only with the

ring-bond matching and valence conditions, and therefore, their decoder sometimes fails to generate valid molecules.

Another approach is to substitute another molecular representation for SMILES so that the output is guaranteed to be valid. As far as we know, only the very recent paper by Jin et al. (2018) takes this approach. They represent a molecule by fragments connected in a tree structure. While their work for the first time reports 100% validity of decoded molecules, their method requires multiple neural networks other than VAE and the predictor. Our work further pushes along this direction by formalizing the tree representation in terms of HRG. This formalization allows us to model atom-level connections between fragments along with the stereochemistry information, and we realize 100% validity using VAE and the predictor only.

## 8. Empirical Studies

We evaluate the effectiveness of MHG in the molecular optimization domain. In particular, we are interested in the case when unlabeled molecules are abundant but the number of function evaluations is limited due to its cost. We basically follow Jin et al.’s experimental procedures<sup>1</sup>, and the baseline results are copied from the existing papers (Jin et al., 2018; You et al., 2018) when appropriate.

**Purposes and Baselines.** As explained in Section 7, there are two complementary approaches to molecular optimization: VAE-based and RL-based ones. The purposes of the empirical studies are to answer the following research questions: **(Q1)** Do RL-based approaches perform better than VAE-based ones when the number of function evaluations is *not* limited (as reported in You et al.’s paper)? **(Q2)** Do VAE-based approaches generally work better than RL-based ones when the number of function evaluations is limited? **(Q3)** Does MHG-VAE outperform the existing VAE-based methods? **(Q2)** and **(Q3)** are our primary concerns. To this end, as baseline methods, we employ CVAE (Gómez-Bombarelli et al., 2018), GVAE (Kusner et al., 2017), SD-VAE (Dai et al., 2018), and JT-VAE (Jin et al., 2018) as VAE-based approaches, and GCPN (You et al., 2018) as an RL-based approach.<sup>2</sup> For their details, see Section 7.

**Dataset.** We use the ZINC dataset following the existing work. This dataset is extracted from the ZINC database (Irwin et al., 2012) and contains 220,011 molecules for training, 24,445 for validation, and 5,000 for testing. For basic statistics of MHG inferred using this dataset, please refer to

<sup>1</sup>Concurrently with our work, some researchers propose to standardize the benchmark (Brown et al., 2019) to ease comparative studies. We would like to test our method on it as future work.

<sup>2</sup>The algorithm by Zhou et al. (2018) is not used because the open-sourced program consumed more than 96GB memory and was unable to reproduce the result in our environment.

Appendix D. For the target chemical property to be maximized, we employ a standardized penalized logP following the existing work:

$$f(m) = \widehat{\log P}(m) - \widehat{SA}(m) - \widehat{\text{cycle}}(m), \quad (4)$$

where logP is the octanol-water partition coefficient, SA is the synthetic accessibility score, and cycle is the size of the longest ring subtracted by six (if its size is less than six, the function returns 0), and the hat represents that the function is standardized using the values calculated on the training set. For example, letting  $\mu_{\log P}$  and  $\sigma_{\log P}$  be the sample mean and standard deviation of logP calculated using the training set,  $\widehat{\log P}(m) = (\log P(m) - \mu_{\log P})/\sigma_{\log P}$ .

### 8.1. Reconstruction Rate

We first investigate the quality of Enc and Dec of the VAE-based methods, which has much influence on the performance of molecular optimization.

**Protocol.** For each molecular graph  $m$  in the test set, we obtain its reconstruction as  $m' = \text{Dec}(\text{Enc}(m))$ . If  $m$  and  $m'$  are isomorphic, we regard the reconstruction succeeds. We repeat the above procedure using all of the test molecules 100 times, and report the mean reconstruction success rate.

To investigate the quality of the latent space, we evaluate the success rate of decoding random latent vectors. We sample  $z$  from  $\mathcal{N}(0, I)$  and decode it to obtain  $m = \text{Dec}(z)$ . If  $m$  is valid, we regard the decode succeeds. We repeat this procedure 1,000 times and report the success rate.

**Result.** According to Table 1 (left), our method clearly improves the reconstruction rate, which justifies our molecular modeling approach.

### 8.2. Global Molecular Optimization

We then investigate the performance on global molecular optimization using two scenarios. The first scenario assumes that the function evaluation cost is negligible and the algorithms can query an arbitrary number of target properties, which is used to answer **(Q1)** and **(Q3)**. This scenario is dubbed as the *unlimited oracle case*. All of the existing studies assume this scenario. The second one assumes that it is expensive and only a limited number of oracle calls are allowed, which is used to answer **(Q2)** and **(Q3)**. This scenario is dubbed as the *limited oracle case*. Since our problem setting assumes the expensive case, the second scenario is of primary interest, and the first scenario is examined for completeness.

**Protocol 1 (unlimited oracle case).** For our method, we first obtain latent representations by Algorithm 1. Then, we apply PCA to the latent vectors to obtain 40-dimensional latent representations. Then, we run Algorithm 2 with  $M =$

Table 1. Reconstruction rate, predictive performance, and global molecular optimization with the unlimited oracle. GCPN, the RL-based method, outperforms VAE-based ones when the target evaluation cost is negligible (though, which is *not* our focus).

Method	% Reconst.	Valid prior	Log likelihood	RMSE	Unlimited oracle case				
					1st	2nd	3rd	50th	Top 50 Avg.
CVAE	44.6%	0.7%	$-1.812 \pm 0.004$	$1.504 \pm 0.006$	1.98	1.42	1.19	–	–
GVAE	53.7%	7.2%	$-1.739 \pm 0.004$	$1.404 \pm 0.006$	2.94	2.89	2.80	–	–
SD-VAE	76.2%	43.5%	$-1.697 \pm 0.015$	$1.366 \pm 0.023$	4.04	3.50	2.96	–	–
JT-VAE	76.7%	100%	$-1.658 \pm 0.023$	$1.290 \pm 0.026$	5.30	4.93	4.49	3.48	3.93
GCPN	–	–	–	–	7.98	7.85	7.80	–	–
<b>Ours</b>	94.8%	100%	$-1.323 \pm 0.003$	$0.959 \pm 0.002$	5.56	5.40	5.34	4.12	4.49

Table 2. Global molecular optimization with the limited oracle. Our method outperforms the others including GCPN.

Method	Limited oracle case				
	1st	2nd	3rd	50th	Top 50 Avg.
JT-VAE	1.69	1.68	1.60	-9.93	-1.33
GCPN	2.77	2.73	2.34	0.91	1.36
<b>Ours</b>	5.24	5.06	4.91	4.25	4.53

50,  $K = 5$ . As a result, we obtain 250 novel molecules. We repeat this procedure ten times, resulting in 2,500 novel molecules. We report the log-likelihood and root mean-squared error (RMSE) of GP evaluated on the test set, top three molecule property scores out of 2,500, and the mean of top 50 target properties. Since BO is a highly random procedure, examining only top three molecules could lead to unfair comparison, and we suggest to examine statistics of top- $K$  molecules.

For the baseline methods, we simply copied the results from You et al.’s paper, where the VAE-based methods use the same protocol as ours. Note that the number of queries by GCPN, the RL-based method, is larger than those of VAE-based methods. In fact, the default configuration of GCPN requires  $5 \times 10^7$  steps.

**Protocol 2 (limited oracle case).** In this scenario, we compare all of the methods under the same number of queries. For our method and JT-VAE, we initialize GP with  $N = 250$  labeled molecules randomly selected from the training set, and run Algorithm 2 with  $M = 1$ ,  $K = 250$ . We use GPpyOpt (The GPpyOpt authors, 2016) as the BO module. For GCPN, we run the algorithm and regard the first 500 molecules as the output. For each method, this procedure is repeated ten times, and obtain 2,500 novel molecules for the VAE-based methods and 5,000 novel molecules for GCPN. We report top three molecule property scores out of the novel molecules and statistics of top 50 target properties.

**Result (the unlimited oracle case).** Table 1 (mid) shows the predictive performance of GP. Our method achieves better scores than the others, indicating that our latent space well encodes features necessary to predict the property.

Table 1 (right) reports the top three target properties as well as the minimum and average scores of top 50 molecules obtained via Protocol 1. Ours achieves the best scores among the VAE-based methods, and GCPN reports molecules with the highest target properties, which answers (Q1) and (Q3) in the affirmative. Furthermore, when we focus on statistics of top 50 molecules, ours achieves better scores than JT-VAE. These results suggest that our method is more likely to discover better molecules than the other VAE-based ones, which also supports us to answer (Q3) in the affirmative.

**Result (limited oracle case).** Table 2 shows the result obtained via Protocol 2. Our method clearly outperforms the other methods including GCPN, answering (Q2) in the affirmative. It is notable that our method in the limited oracle case performs almost comparably to our method in the unlimited oracle case. This result supports the effectiveness of our method especially in the limited oracle case.

## 9. Conclusion and Future Work

We have developed the molecular hypergraph grammar variational autoencoder (MHG-VAE). Our key idea is to employ MHG to represent a molecular graph as a parse tree, which is fed into VAE. Since MHG models the atom-level connections as well as the stereochemistry information, MHG-VAE can learn a pair of encoder and decoder using a single VAE. The highlights of our experiments include (i) MHG-VAE achieves the best performance among VAE-based methods and (ii) MHG-VAE performs better than the state-of-the-art RL-based method called GCPN when the number of target function evaluations is limited.

A future research direction will be to optimize MHG with respect to some goodness criteria such as the minimum description length (Jonker et al., 2004) or a Bayesian criterion (Chen, 1995). Since the resultant MHG depends on tree decompositions, we can optimize MHG by manipulating tree decompositions.

Another direction involves retrosynthetic analysis, which derives a pathway to synthesize the target molecule. With this capability, we will be able to immediately examine the property by synthesizing the output molecule.

## Acknowledgments

This work was supported by JST CREST Grant Number JPMJCR1304, Japan, and JSPS KAKENHI Grant Numbers 15H05711, Japan. The author would like to thank Dr. Masakazu Ishihata for his helpful discussion.

## References

- Aguiñaga, S., Palacios, R., Chiang, D., and Weninger, T. Growing graphs from hyperedge replacement graph grammars. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, pp. 469–478, 2016.
- Brown, N., Fiscato, M., Segler, M. H., and Vaucher, A. C. GuacaMol: Benchmarking models for de novo molecular design. *Journal of Chemical Information and Modeling*, 59(3):1096–1108, 2019.
- Chen, S. F. Bayesian grammar induction for language modeling. In *Proceedings of the 33rd Annual Meeting on Association for Computational Linguistics*, pp. 228–235, 1995.
- Dai, H., Tian, Y., Dai, B., Skiena, S., and Song, L. Syntax-directed variational autoencoder for structured data. In *Proceedings of the Sixth International Conference on Learning Representations*, 2018.
- Drewes, F., Kreowski, H.-J., and Habel, A. *Hyperedge replacement graph grammars*, volume 1, chapter 2, pp. 95–162. 1997.
- Gómez-Bombarelli, R., Aguilera-Iparraguirre, J., Hirzel, T. D., Duvenaud, D., Maclaurin, D., Blood-Forsythe, M. A., Chae, H. S., Einzinger, M., Ha, D.-G., Wu, T., Markopoulos, G., Jeon, S., Kang, H., Miyazaki, H., Numata, M., Kim, S., Huang, W., Hong, S. I., Baldo, M., Adams, R. P., and Aspuru-Guzik, A. Design of efficient molecular organic light-emitting diodes by a high-throughput virtual screening and experimental approach. *Nature Materials*, 15(10):1120–1127, 2016.
- Gómez-Bombarelli, R., Wei, J. N., Duvenaud, D., Hernández-Lobato, J. M., Sánchez-Lengeling, B., Sheberla, D., Aguilera-Iparraguirre, J., Hirzel, T. D., Adams, R. P., and Aspuru-Guzik, A. Automatic chemical design using a data-driven continuous representation of molecules. *ACS Central Science*, 2018.
- Guimaraes, G. L., Sanchez-Lengeling, B., Outeiral, C., Farias, P. L. C., and Aspuru-Guzik, A. Objective-reinforced generative adversarial networks (ORGAN) for sequence generation models. Technical report, arXiv1705.10843, 2017.
- Irwin, J. J., Sterling, T., Mysinger, M. M., Bolstad, E. S., and Coleman, R. G. ZINC: A free tool to discover chemistry for biology. *Journal of Chemical Information and Modeling*, 52(7):1757–1768, 2012.
- Jin, W., Barzilay, R., and Jaakkola, T. Junction tree variational autoencoder for molecular graph generation. In *Proceedings of the Thirty-fifth International Conference on Machine Learning*, 2018.
- Jonyer, I., Holder, L. B., and Cook, D. J. MDL-based context-free graph grammar induction and applications. *International Journal on Artificial Intelligence Tools*, 13(1):65–79, 2004.
- Jorgensen, W. L. Efficient drug lead discovery and optimization. *Accounts of Chemical Research*, 42(6):724–733, 2009.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. In *Proceedings of the International Conference on Learning Representations*, 2014.
- Kusner, M. J., Paige, B., and Hernández-Lobato, J. M. Grammar variational autoencoder. In *Proceedings of the 35th International Conference on Machine Learning*, 2017.
- Močkus, J. On Bayesian methods for seeking the extremum. In Marchuk, G. I. (ed.), *Optimization Techniques IFIP Technical Conference Novosibirsk, July 1–7, 1974*, pp. 400–404, 1975.
- The GPyOpt authors. GPyOpt: A Bayesian optimization framework in Python. <http://github.com/SheffieldML/GPyOpt>, 2016.
- Weninger, D. SMILES, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of Chemical Information and Computer Sciences*, 28(1):31–36, 1988.
- You, J., Liu, B., Ying, Z., Pande, V., and Leskovec, J. Graph convolutional policy network for goal-directed molecular graph generation. In *Advances in Neural Information Processing Systems 31*, pp. 6412–6422, 2018.
- Zhou, Z., Kearnes, S., Li, L., Zare, R. N., and Riley, P. Optimization of molecules via deep reinforcement learning. Technical report, arXiv1810.08678, 2018.