

---

# A Graph to Graphs Framework for Retrosynthesis Prediction

---

Chence Shi<sup>1</sup> Minkai Xu<sup>2</sup> Hongyu Guo<sup>3</sup> Ming Zhang<sup>1</sup> Jian Tang<sup>4,5,6</sup>

## Abstract

A fundamental problem in computational chemistry is to find a set of reactants to synthesize a target molecule, a.k.a. retrosynthesis prediction. Existing state-of-the-art methods rely on matching the target molecule with a large set of reaction templates, which are very computationally expensive and also suffer from the problem of coverage. In this paper, we propose a novel template-free approach called G2Gs by transforming a target molecular graph into a set of reactant molecular graphs. G2Gs first splits the target molecular graph into a set of synthons by identifying the reaction centers, and then translates the synthons to the final reactant graphs via a variational graph translation framework. Experimental results show that G2Gs significantly outperforms existing template-free approaches by up to 63% in terms of the top-1 accuracy and achieves a performance close to that of state-of-the-art template-based approaches, but does not require domain knowledge and is much more scalable.

## 1. Introduction

Retrosynthesis, which devotes to find a set of reactants to synthesize a target molecule, is of crucial importance to the synthesis planning and drug discovery. The problem is challenging as the search space of all possible transformations is huge by nature. For decades, people have been seeking to aid retrosynthesis analysis with modern computer techniques (Corey & Wipke, 1969). Among them, machine learning plays a vital role and significant progress has been made recently (Szymkuć et al., 2016; Coley et al., 2018).

Existing machine learning works on retrosynthesis predic-

tion mainly fall into two categories: template-based (Coley et al., 2017b; Segler & Waller, 2017; Dai et al., 2019) and template-free models (Liu et al., 2017; Karpov et al., 2019). The template-based approaches match the target molecule with a large set of reaction templates, which define the subgraph patterns of a set of chemical reactions. For example, (Coley et al., 2017b) proposed a similarity-based approach to select reaction templates for the target molecule. (Segler & Waller, 2017; Baylon et al., 2019) cast rule selection as a multi-class classification problem. Recently, (Dai et al., 2019) treats chemistry knowledge as logic rules and directly models the joint probability of rules and reactants, which achieves the new state of the art. Despite their great potential for synthesis planning, template-based methods, however, not only require expensive computation but also suffer from poor generalization on new target structures and reaction types. Another line of research for retrosynthesis prediction (Liu et al., 2017; Karpov et al., 2019) bypasses reaction templates and formulates retrosynthesis prediction as a sequence-to-sequence problem. Such approaches leverage the recent advances in neural machine translation (Bahdanau et al., 2014; Vaswani et al., 2017) and the SMILES (Weininger, 1988) representation of molecules. However, SMILES representation assumes a sequential order between the atoms in a molecule, which cannot effectively reflect the complex relationships between atoms in a molecule. As a result, these approaches fail to capture the rich chemical contexts and their interplays of molecules, resulting in unsatisfactory predictive performance.

To address the aforementioned issues, in this paper we represent each molecule as a graph and formulate retrosynthesis prediction as a graph-to-graphs translation problem. The so-called G2Gs leverages the powerful representation of graph for molecule and is a novel template-free approach, which is trained with an extensive collection of molecule reaction data. Specifically, it consists of two key components: (1) a reaction center identification module, which splits synthons from the target molecule and reduces the one-to-many graph translation problem into multiple one-to-one translation processes; (2) a variational graph translation module, which translates a synthon to a final reactant graph. We aim to model the probability of reactant graph  $G$  conditioned on the synthon  $S$ , i.e.,  $P(G|S)$ . As a synthon could be potentially translated to different reactants in different reaction contexts,

<sup>1</sup>Department of Computer Science, School of EECS, Peking University <sup>2</sup>Shanghai Jiao Tong University <sup>3</sup>National Research Council Canada <sup>4</sup>Montréal Institute for Learning Algorithms (MILA) <sup>5</sup>Canadian Institute for Advanced Research (CIFAR) <sup>6</sup>HEC Montréal. Correspondence to: Chence Shi <chenceshi@pku.edu.cn>, Jian Tang <jian.tang@hec.ca>.

a latent code  $z$  is introduced to handle the uncertainty of reactant prediction, *i.e.*,  $P(G|z, S)$ . Following existing work on graph generation (You et al., 2018a; Liu et al., 2018; Shi et al., 2020), we formulate reactant graph generation as a sequential decision process, more specifically, sequentially generating nodes and edges. The whole graph translation process can be efficiently and effectively optimized with the variational inference approach (Kingma & Welling, 2014).

We evaluate our model on the benchmark data set USPTO-50k derived from a patent database (Lowe, 2012), and compare it with both template-based and template-free approaches. We experimentally show that G2Gs significantly outperforms existing template-free baselines by up to 63% in terms of the top-1 accuracy. These numbers are also approaching those obtained by the state-of-the-art template-based strategies, but our method excludes the need for domain knowledge and scales well to larger data sets, making it particularly attractive in practice.

## 2. Related Work

**Retrosynthesis Prediction** Prior works on retrosynthesis prediction are primarily based on reaction templates, which are either hand-crafted by human experts (Hartenfeller et al., 2011; Szymkuć et al., 2016) or extracted from large chemical databases automatically (Coley et al., 2017a). Since there are hundreds of qualified templates for a target molecule by subgraph matching, selecting templates that lead to chemically feasible reactions remains a crucial challenge for these approaches. To cope with such challenge, (Coley et al., 2017b) proposed to select templates based on similar reactions in the data set. (Segler & Waller, 2017; Baylon et al., 2019) further employed neural models for rule selection. The state-of-the-art method (Dai et al., 2019) leveraged the idea that the templates and reactants are hierarchically sampled from their conditional joint distributions. Such template-based approaches, however, still suffer from poor generalization on unseen structures and reaction types, and the computational cost for subgraph isomorphism in these methods is often prohibitively expensive.

To overcome the limitation of template-based methods, template-free approaches have recently been actively investigated. (Liu et al., 2017; Karpov et al., 2019) formulated the task as a sequence-to-sequence problem on SMILES representation of molecules, and the same idea was adopted in (Schwaller et al., 2017; 2018) for its dual problem, *i.e.*, organic reaction prediction. However, these approaches are apt to ignore the rich chemical contexts contained in the graph structures of molecules, and the validity of the generated SMILES strings can not be ensured, resulting in unsatisfactory performance. In contrast, our G2Gs framework, directly operating on the graph structures of molecules, is able to generate 100% chemically valid predictions with

high accuracy, while excludes the need for reaction templates and computationally expensive graph isomorphism.

Our work is also related to the chemical reaction prediction method presented in (Jin et al., 2017) in terms of learning a neural model for reaction center identification. Nevertheless, both the definition of the reaction center and the targeted task, namely retrosynthesis prediction, distinguish our strategy from their algorithm.

**Molecular Graph Generation** Various deep generative models for molecular graph generation have recently been introduced (Segler et al., 2017; Olivecrona et al., 2017; Samanta et al., 2018; Neil et al., 2018; You et al., 2018a; Liu et al., 2018; Shi et al., 2020). Our G2Gs framework is closely related to the state-of-the-art models that decompose the generation process into a sequence of graph transformation steps, including (You et al., 2018a; Liu et al., 2018; Shi et al., 2020). In these approaches, the generation procedure is formulated as a sequential decision making process by dynamically adding new nodes and edges based on current subgraph structures. For example, (You et al., 2018a) introduced a reinforce policy network for the decision making. (Shi et al., 2020) presented an alternative sequential generation algorithm based on autoregressive flow called GraphAF. (Liu et al., 2018) combined sequential generation with a variational autoencoder to enable continuous graph optimization in the latent space.

Unlike the aforementioned approaches, we here leverage graph generation for retrosynthesis prediction, where novel components have to be devised for this specific task. For example, in the retrosynthesis scenario, given a specified product molecule, the reactants could be slightly different with diverse reaction conditions such as reagent and temperature. To cope with such intrinsic multi-modality problem, in G2Gs we also explicitly introduce low-dimensional latent variables to modulate the sequential graph translation process, aiming at capturing the diverse output distributions.

## 3. The Graph to Graphs Framework

In this paper, we formulate the retrosynthesis task as a one-to-many graph-to-graphs translation problem. In specific, we first employ a graph neural network to estimate the reactivity score of all atom pairs of the product graph, and the atom pair with the highest reactivity score above a threshold will be selected as the reaction center. We then split the product graph into synthons by disconnecting the bonds of the reaction center resulted. Finally, basing on the obtained synthons, the reactants are generated via a series of graph transformations, where a latent vector is employed to encourage the model to capture the transformation uncertainty and generate diverse predictions. The proposed framework

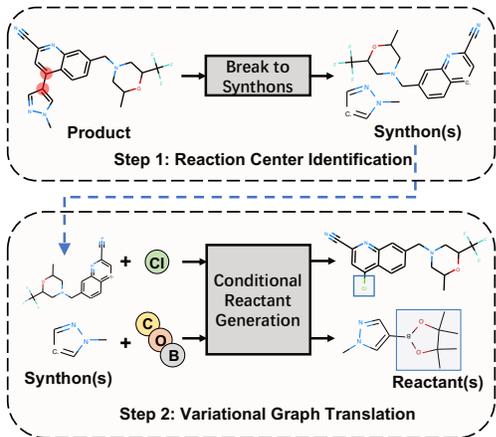


Figure 1. The overall framework of the proposed method. The reaction center identified by G2Gs is marked in red. The product graph is first split into synthons by disconnecting the reaction center. Based on resulted synthons, reactants are then generated via a series of graph transformations. The generated molecule scaffolds are bounded by blue bounding box.

is illustrated in Figure 1, and will be discussed in detail next.

### 3.1. Preliminaries

**Notation** In this paper, a molecule is represented as a labeled graph  $G = (A, X)$ , where  $A$  is the adjacency matrix and  $X$  the matrix of node features. Let the number of atoms be  $n$ , the number of bond types be  $b$ , and the dimension of node features be  $d$ , then we have  $A \in \{0, 1\}^{n \times n \times b}$  and  $X \in \{0, 1\}^{n \times d}$ .  $A_{ijk} = 1$  here indicates a bond with type  $k$  between the  $i^{th}$  and  $j^{th}$  nodes.

**Retrosynthesis Prediction** Formally, a chemical reaction can be represented as a pair of two sets of molecules  $(\{G_i\}_{i=1}^{N_1}, \{G_j\}_{j=1}^{N_2})$ , where  $G_i$  denotes a reactant graph and  $G_j$  a product graph. In retrosynthesis prediction, given the set of products  $\{G_j\}_{j=1}^{N_2}$ , the goal is to precisely predict the set of reactants  $\{G_i\}_{i=1}^{N_1}$ . Following existing work, in this paper we focus on the standard single-outcome reaction case, *i.e.*,  $N_2 = 1$ , and thus simplifying the notation of a reaction  $r$  as  $(\{G_i\}_{i=1}^{N_1}, G_p)$ .

**Reaction Center and Synthon** Unlike that in (Jin et al., 2017), the phrase *reaction center* here is used to represent an atom pair  $(i, j)$  that satisfies two conditions: (1) there is a bond between the  $i^{th}$  and  $j^{th}$  nodes in the product graph; (2) there is no bond between the  $i^{th}$  and  $j^{th}$  nodes in the reactant graph. Also, synthons are subgraphs extracted from the products by breaking the bonds in the reaction centers. These synthons can later be transformed into reactants, and a synthon may not be a valid molecule.

**Molecular Graph Representation Learning** Graph Convolutional Networks (GCN) (Duvenaud et al., 2015; Kearnes

et al., 2016; Kipf & Welling, 2016; Gilmer et al., 2017; Schütt et al., 2017) have achieved great success in representation learning for computational chemistry. Given a molecular graph  $G = (A, X)$ , we adopt a variant of Relational GCN (R-GCN) (Schlichtkrull et al., 2018) to compute both the node embeddings and the graph-level embedding. Let  $k \in \mathbb{R}$  be the embedding dimension and  $H^l \in \mathbb{R}^{n \times k}$  the node embeddings at the  $l^{th}$  layer computed by the R-GCN ( $H^0 = X$ ).  $H_i^l$  represents the embedding of the  $i^{th}$  atom. At each layer, the R-GCN calculates the embedding for each node by aggregating messages from different edge types:

$$H^l = \text{Agg}(\text{ReLU}(\{E_i H^{l-1} W_i^l\} | i \in (1, \dots, b))) \quad (1)$$

where  $E_i = A_{[:, :, i]} + I$  denotes the adjacency matrix of the  $i^{th}$  edge type with self-loop and  $W_i^l$  is the trainable weight matrix for the  $i^{th}$  edge type.  $\text{Agg}(\cdot)$  denotes an aggregation function chosen from summation, averaging and concatenation. We stack  $L$  R-GCN layers to compute the final node embeddings  $H^L$ . The entire graph-level embedding  $h_G$  can also be calculated by applying a Readout( $\cdot$ ) function to  $H^L$  (Hamilton et al., 2017), *e.g.*, summation.

### 3.2. Reaction Center Identification

Given a chemical reaction  $(\{G_i\}_{i=1}^{N_1}, G_p)$ , we first derive a binary label matrix  $Y \in \{0, 1\}^{n \times n}$  for each atom pair (*i.e.*, bond) in the product  $G_p$ , indicating the reaction centers as defined in Section 3.1;  $Y_{ij} = 1$  here indicates that the bond between  $i^{th}$  and  $j^{th}$  node in  $G_p$  is a reaction center. It is worth noting that, both the reactants and the product are *atom-indexed*. That is, each atom in the reactant set is associated with a unique index. This property enables us to identify the reaction centers by simply comparing each pair of atoms in the product  $G_p$  with that in a reactant  $G_i$ , forming the binary label matrix  $Y$ . With such label matrix, the Reaction Center Identification procedure in G2Gs is then formulated as a binary link prediction task as follows.

We use a  $L$ -layer R-GCN defined in Section 3.1 to compute the node embeddings  $H^L$  and the entire graph embedding  $h_{G_p}$  of product  $G_p$ :

$$H^L = \text{R-GCN}(G_p), h_{G_p} = \text{Readout}(H^L). \quad (2)$$

To estimate the reactivity score of the atom pair  $(i, j)$ , the edge embedding  $e_{ij}$  is formed by concatenating the node embeddings of the  $i^{th}$  and  $j^{th}$  nodes as well as the one-hot bond type feature. In addition, since a  $L$ -layer R-GCN can only gather information within  $L$  hops of the center node, but the reactivity of a reaction center may also be affected by the remote atoms, we also enable the edge embedding  $e_{ij}$  to take into account the graph embedding (*i.e.*, global structure), so as to leverage the knowledge of the remote atoms. Formally, we have:

$$e_{ij} = H_i^L \parallel H_j^L \parallel A_{ij} \parallel h_{G_p} \quad (3)$$

where  $\parallel$  denotes the vector concatenation operation.  $H_i^L \in \mathbb{R}^k$  denotes the  $i^{\text{th}}$  row of  $H^L$  and  $A_{ij} = A_{[i,j,:]} \in \{0,1\}^b$  is the edge type of the atom pair  $(i,j)$  in  $G_p$ . The final reactivity score of the atom pair  $(i,j)$  is calculated as:

$$s_{ij} = \sigma(m_r(e_{ij})) \quad (4)$$

where  $m_r$  is a feedforward network that maps  $e_{ij}$  to a scalar, and  $\sigma(\cdot)$  denotes the Sigmoid function. In the case that a certain reaction type is known during retrosynthesis analysis, we can also include this information by concatenating its embedding to the input of the feedforward network, *i.e.*,  $e_{ij}$ .

For learning, the reaction center identification module can be optimized by maximizing the cross entropy of the binary label matrix  $Y$  as follows:

$$\mathcal{L}_1 = - \sum_r \sum_{i \neq j} \lambda Y_{ij} \log(s_{ij}) + (1 - Y_{ij}) \log(1 - s_{ij}) \quad (5)$$

where  $\lambda \in [1, +\infty)$  is a weighting hyper-parameter, with the following purpose. In practice, among all the atom pairs there are typically a few reaction centers. As a result, the output logit of the feedforward network  $m_r$  is often extremely small. Through the weighting  $\lambda$ , we are able to alleviate such issue caused by the imbalanced class distributions.

During inference, we first calculate the reactivity scores of all atom pairs, and then select the highest one above a threshold as the reaction center. Alternatively, one can select the top- $k$  atom pairs with the highest scores above a threshold as the reaction center. This scenario will yield more diverse synthesis routes but with the cost of inference time.

After collecting reaction centers, we then disconnect the bonds of the reaction centers in  $G_p$ , and treat each connected subgraph in  $G_p$  as a synthon. Note that, in the case that none of the reaction center is identified, the product  $G_p$  itself will be considered as a synthon. Doing so, we can extract all the synthons from the product  $G_p$ , and then formulate a one-to-one graph translation procedure to translate each resulted synthon to a final reactant graph. We will discuss this translation process in detail next.

### 3.3. Reactants Generation via Variational Graph Translation

Given a chemical reaction  $(\{G_i\}_{i=1}^{N_1}, G_p)$ , we denote the set of synthons extracted from  $G_p$  (as discussed in Section 3.2) as  $\{S_i\}_{i=1}^{N_1}$ . For simplicity, we omit the subscript and denote a translation pair as  $(S, G)$ . In this setting, our goal is to learn a conditional generative model  $p(G|S)$  that recovers the reactant molecule domain ( $G$ ) from the synthon molecule domain ( $S$ ). It is worth noting that, an intrinsic issue is that the same synthon can be translated to different reactants, which is known as the multi-modality problem. In

order to mitigate such multi-modality problem and encourage the module to model the diverse distribution of reactants, we incorporate a low-dimensional latent vector  $z$  to capture the uncertainty for the graph translation process. Details are discussed next.

#### 3.3.1. THE GENERATIVE MODEL

We build upon previous works on graph generation models (Li et al., 2018; You et al., 2018b). The generation of graph  $G$  is conditioned on both the  $S$  and the latent vector  $z$ . In detail, we first autoregressively generate a sequence of graph transformation actions  $(a_1, \dots, a_T)$ , and then apply them on the initial synthon graph  $S$ . Here  $a_t$  is a one-step graph transformation (*i.e.*, action) acting as a modification to the graph. Formally, let  $\mathcal{T}$  be the collection of all trajectories  $(a_1, \dots, a_T)$  that can translate synthons  $S$  to target reactants  $G$  and  $t \in \mathcal{T}$  be a possible trace. Then, we factorize modeling  $p(G|z, S)$  into modeling the joint distribution over the sequence of graph transformation actions  $p(t|z, S)$ . The connection between  $p(G|z, S)$  and  $p(t|z, S)$  is illustrated in Section 3.3.2. With such a generative model, a synthon can be translated into a reactant by sampling action sequences from the distribution. Next, we describe the details of the generative procedure.

Let  $S^i$  denote the graph after applying the sequence of actions  $a_{1:i}$  to  $S$ , and  $S^0 = S$ . Then we have  $p(S^i|S^{i-1}, z) = p(a_i|S^{i-1}, z)$ . In previous graph generation models (Li et al., 2018; You et al., 2018b), each decision  $a_i$  is conditioned on a full history of the generation sequence  $(S^0, \dots, S^{i-1})$ , which leads to the stability and scalability problems arising from the optimization process. To alleviate these issues, our graph translation dynamics leverages the assumption of a Markov Decision Process (MDP), which satisfies the Markov property that  $p(S^i|S^{i-1}, z) = p(S^i|S^{i-1}, \dots, S^0, z)$ . The MDP formulation means that each action is only conditioned on the graph that has been modified so far. Hence the graph translation model  $p(t|z, S)$  can be naturally factorized as follows:

$$p(t|z, S) = p(a_{1:T}|z, S) = \prod_{i=1}^T p(a_i|z, S^{i-1}). \quad (6)$$

Before we detail the parameterization of distribution  $P(a_i|z, S^{i-1})$ , we formally introduce the definition of an action. An action  $a_i$  is a tuple with four elements:

$$a_i = (a_i^1, a_i^2, a_i^3, a_i^4). \quad (7)$$

Assuming the number of atom types is  $m$ , then  $a_i^1 \in \{0, 1\}^2$  predicts the termination of the graph translation procedure;  $a_i^2 \in \{0, 1\}^n$  indicates the first node to focus;  $a_i^3 \in \{0, 1\}^{n+m}$  indicates the second node to focus;  $a_i^4 \in \{0, 1\}^b$  predicts the type of bond between two

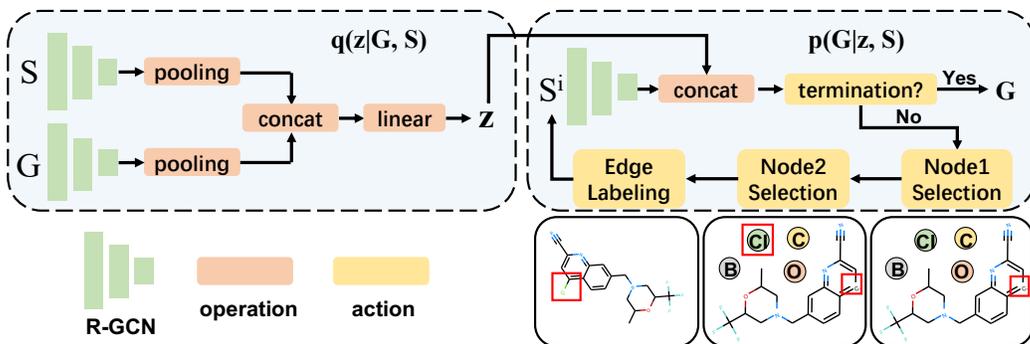


Figure 2. Illustration of the proposed variational graph translation module, including the encoder  $q(z|G, S)$  (left) and the generative model  $p(G|z, S)$  (right). The phases of the generation procedure are shown in the generative model. With a sampled latent vector  $z$ , the synthon graph  $S$  enters a loop with nodes selection and edge labeling until the termination state is predicted.

nodes. Then the distribution  $p(a_i|z, S^{i-1})$  can be further decomposed into three parts: (1) termination prediction  $p(a_i^1|z, S^{i-1})$ ; (2) nodes selection  $p(a_i^{2:3}|z, S^{i-1}, a_i^1)$ ; (3) edge labeling  $p(a_i^4|z, S^{i-1}, a_i^{1:3})$ . We will discuss the parameterization of these components in detail next.

**Termination Prediction** Denote a  $L$ -layer R-GCN (Section 3.1) that can compute the node embeddings of an input graph as  $\mathcal{R}(\cdot)$ . We parameterize the distribution as:

$$H = \mathcal{R}(S^{i-1}), h_S = \text{Readout}(H) \\ p(a_i^1|z, S^{i-1}) = \tau(m_t(h_S, z)) \quad (8)$$

where  $\tau(\cdot)$  denotes the softmax function, and  $m_t(\cdot)$  is a feedforward network. At each transformation step, we sample  $a_i^1 \sim p(a_i^1|z, S^{i-1})$ . If  $a_i^1$  indicates the termination, then the graph translation process will stop, and the current graph  $S^{i-1}$  is treated as the final reactant  $G$  generated by the module.

**Nodes Selection** Let the set of possible atoms (*e.g.*, carbon, oxygen) to be added during graph translation be  $\{v_1, \dots, v_m\}$  and denote the collection as  $V = \bigcup_{i=1}^m v_i$ . We first extend the graph  $S^{i-1}$  to  $\tilde{S}^{i-1}$  by adding isolated atoms, *i.e.*,  $\tilde{S}^{i-1} = S^{i-1} \cup V$ . The first node is selected from atoms in  $S^{i-1}$ , while the second node is selected from  $\tilde{S}^{i-1}$  conditioned on the first node, by concatenating its embeddings with embeddings of each atom in  $\tilde{S}^{i-1}$ :

$$p(a_i^2|z, S^{i-1}, a_i^1) = \tau(\beta_1 \odot m_f(\mathcal{R}(\tilde{S}^{i-1}), z)) \\ a_i^2 \sim p(a_i^2|z, S^{i-1}, a_i^1) \\ p(a_i^3|z, S^{i-1}, a_i^{1:2}) = \tau(\beta_2 \odot m_s(\mathcal{R}(\tilde{S}^{i-1}), z, a_i^2)) \\ a_i^3 \sim p(a_i^3|z, S^{i-1}, a_i^{1:2}) \quad (9)$$

where  $m_f(\cdot)$  and  $m_s(\cdot)$  are feedforward networks.  $\beta_1$  and  $\beta_2$  are masks to zero out the probability of certain atoms being selected. Specifically,  $\beta_1$  forces the module to select node from  $S^{i-1}$ , and  $\beta_2$  prevents the first node from being selected again. In this case, only the second node

can be selected from  $V$ , and it corresponds to adding a new atom to  $S^{i-1}$ . The distribution of node selection  $p(a_i^{2:3}|z, S^{i-1}, a_i^1)$  is the product of the two conditional distributions described above.

**Edge Labeling** The distribution for edge labeling is parameterized as follows:

$$p(a_i^4|z, S^{i-1}, a_i^{1:3}) = \tau(m_e(\mathcal{R}(\tilde{S}^{i-1}), z, a_i^{2:3})) \\ a_i^4 \sim P(a_i^4|z, S^{i-1}, a_i^{1:3}) \quad (10)$$

where  $m_e(\cdot)$  is a feedforward networks. The knowledge of reaction classes can be incorporated in the same way as in Section 3.2.

Given these distributions, the distribution  $p(t|z, S)$  can be parameterized according to eqs. (6) and (8) to (10). Finally, the probability of translating the synthon  $S$  to the final reactant  $G$  can be computed by enumerating all possible graph transformation sequences that translate  $S$  to  $G$ :

$$P(G|z, S) = \sum_{t \in \mathcal{T}} P(t|z, S) \quad (11)$$

### 3.3.2. LEARNING

To learn the parameters of our variational graph translation module, we aim to maximize the log likelihood of the observed translation pair, *i.e.*,  $\log P(G|S)$ . Directly optimizing the objective involves marginalizing the latent variable  $z$ , which is computationally intractable. To this end, we turn to the standard amortized variational inference (Kingma & Welling, 2014) by introducing an approximate posterior  $q(z|G, S)$ , which is modeled as a Gaussian distribution to allow effectively sampling  $z$  via the reparameterization trick. In specific, the mean and the log variance of  $q(z|G, S)$  are parameterized as follows:

$$\mu = m_\mu(h_G \| h_S) \\ \log \sigma^2 = m_\sigma(h_G \| h_S) \\ q(z|G, S) = \mathcal{N}(z | \mu, \text{diag}(\sigma^2)) \quad (12)$$

where the  $h_G$  and  $h_S$  are graph embeddings of  $G$  and  $S$  respectively, computed by the same R-GCN.  $m_\mu(\cdot)$  and  $m_\sigma(\cdot)$  are feedforward networks. The evidence lower bound (ELBO) is then defined as:

$$\mathcal{L}_{\text{ELBO}} = \mathbb{E}_{z \sim q}[\log P(G|z, S)] - \text{KL}[q(z|G, S) \| p(z|S)] \quad (13)$$

where  $\text{KL}[q(\cdot) \| p(\cdot)]$  is the Kullback-Leibler divergence between  $q(\cdot)$  and  $p(\cdot)$ . We here take the prior  $p(z|S)$  as a standard Gaussian  $\mathcal{N}(z|0, I)$ .

**Efficient Training** The computation of  $\log P(G|z, S)$  (eq. 11) is expensive as it requires the summation over all possible graph transformation sequences that translate  $S$  to  $G$ . Here we introduce two strategies that perform well empirically. We first show that  $\log P(G|z, S)$  is lower bounded by the expected log likelihood of all the trajectories  $t = (a_1, \dots, a_T)$  that translate  $S$  to  $G$  using Jensen’s inequality:

$$\begin{aligned} \log P(G|z, S) &= \log \sum_{t \in \mathcal{T}} P(t|z, S) \\ &\geq \log |t| + \frac{1}{|t|} \sum_{t \in \mathcal{T}} \log P(t|z, S) \end{aligned} \quad (14)$$

where  $|t|$  is the number of different action traces that translate  $S$  to  $G$ . In practice, we can throw the constant and evaluate the expectation using Monte Carlo estimation. We further adopt the breadth-first-search (BFS) node-ordering, a widely-used technique in sequential graph generation (You et al., 2018b; Popova et al., 2019; Shi et al., 2020), to reduce the number of valid transformation traces during sampling.

### 3.3.3. GENERATION

To generate a reactant graph, a natural way is to first sample a latent vector  $z$  from the prior  $p(z|S)$ , then sample a trace of graph transformations from  $p(t|z, S)$  (Section 3.3.1), and finally apply these transformations to the synthon  $S$ . However, in our proposed generative model, the probability of invalid actions will be non-zero even if the model is well-trained. As a result, any reactant molecules including invalid ones can be generated if sampling is arbitrarily long. Besides, this process also suffers from the non-trivial exposure bias problem (Bengio et al., 2015). To overcome the above obstacles during sampling, we design a beam search sampling process as follows.

Consider a beam search with size of  $k$ . For the graph generation in the  $i^{\text{th}}$  step, we maintain a candidate set  $\mathcal{S} = \{S^{i,j}\}_{j=1}^k$  with size  $k$ . At the  $i^{\text{th}}$  transformation step, we first calculate the probabilities of all possible actions and sort them, and then select the top  $k$  ranked valid actions for each candidate graph  $S^{i-1,j}$  in  $\mathcal{S}$ . Once this is done for  $k$  graphs in  $\mathcal{S}$ , the top  $k$  graphs among all the generated  $k^2$  graphs are then selected as the candidates for the next

$(i+1)^{\text{th}}$  transformation step. During this beam search, a translation branch will stop if  $i$  reaches the predefined maximum transformation step or  $a_i^1$  indicates a termination. In this scenario, the current graph will be added into a set  $\mathcal{G}$ , and the whole beam search terminates once all translation branches stop. When the beam search finishes, the top  $k$  graphs in  $\mathcal{G}$ , ranked by their likelihoods, will be collected as the final predicted graphs.

### 3.4. Scalability Analysis

Both the Reaction Center Identification and the Variational Graph Translation modules in our G2Gs framework bypass the deployment of reaction templates and take advantage of the representation power of the molecular graph embeddings. As a result, the model size of G2Gs scales linearly *w.r.t* the maximum number of atoms in the molecules and is invariant to the quantity of rules and reactions in the given data sets, making it highly scalable to larger data sets.

## 4. Empirical Studies

### 4.1. Experiment Setup

**Data** We evaluate our approach on the widely used benchmark data set USPTO-50k, which contains 50k atom-mapped reactions with 10 reaction types. Following (Liu et al., 2017), we randomly select 80% of the reactions as training set and divide the rest into validation and test sets with equal size.

**Baselines** We evaluate our strategy using five comparison baselines, including two template-free and three template-based ones. In specific, **Seq2seq** (Liu et al., 2017) is a template-free approach that learns a LSTM (Hochreiter & Schmidhuber, 1997) to translate the SMILES strings of target molecules to reactants. **Transformer** (Karpov et al., 2019) is also a neural sequence-to-sequence model, but it leverages the learning power of the Transformer (Vaswani et al., 2017) for better sequential modeling. **Retrosim** (Coley et al., 2017b) is a data-driven method that selects template for target molecules based on similar reactions in the data set. **Neuralsym** (Segler & Waller, 2017) employs a neural model to rank templates for target molecules. **GLN** (Dai et al., 2019) is the state-of-the-art method, which samples templates and reactants jointly from the distribution learned by a conditional graphical model.

**Evaluation Metrics** Following the existing works (Liu et al., 2017; Dai et al., 2019), we use the top- $k$  exact match accuracy as our evaluation metrics. For comparison purpose, we used  $k$  with 1, 3, 5, and 10 in our experiments. Also, the accuracy was computed by matching the canonical SMILES strings of the predicted molecules with the ground truth.

**Implementation Details** G2Gs is implemented in Py-

torch (Paszke et al., 2017). We use the open-source chemical software RDkit (Landrum, 2016) to preprocess molecules for the training and generate canonical SMILES strings for the evaluation. The R-GCN in G2Gs is implemented with 4 layers and the embedding size is set as 512 for both modules. We use latent codes of dimension  $|z| = 10$ . We train our G2Gs for 100 epochs with a batch size of 128 and a learning rate of 0.0001 with Adam (Kingma & Ba, 2014) optimizer on a single GTX 1080Ti GPU card. The  $\lambda$  is set as 20 for reaction center identification module, and the beam size is 10 during inference. The maximal number of transformation steps is set as 20. We heuristically selected these values based on the validation data set.

Table 1. Top- $k$  exact match accuracy when reaction class is given. Results of all baselines are directly taken from (Dai et al., 2019).

Methods	Top- $k$ accuracy %			
	1	3	5	10
Template-free				
Seq2seq	37.4	52.4	57.0	61.7
G2Gs	<b>61.0</b>	<b>81.3</b>	<b>86.0</b>	<b>88.7</b>
Template-based				
Retrosim	52.9	73.8	81.2	88.1
Neuralsym	55.3	76.0	81.4	85.1
GLN	<b>64.2</b>	<b>79.1</b>	<b>85.2</b>	<b>90.0</b>

Table 2. Top- $k$  exact match accuracy when reaction class is unknown. Results of all baselines are taken from (Dai et al., 2019).

Methods	Top- $k$ accuracy %			
	1	3	5	10
Template-free				
Transformer	37.9	57.3	62.7	/
G2Gs	<b>48.9</b>	<b>67.6</b>	<b>72.5</b>	<b>75.5</b>
Template-based				
Retrosim	37.3	54.7	63.3	74.1
Neuralsym	44.4	65.3	72.4	78.9
GLN	<b>52.5</b>	<b>69.0</b>	<b>75.6</b>	<b>83.7</b>

## 4.2. Predictive Performance

We evaluate the top- $k$  exact match accuracy of the proposed approach in both reaction class known and reaction class unknown settings, with results presented in Tables 1 and 2, respectively. The sets of reactant molecules are generated via beam search (Section 3.3.3) and ranked by their log likelihoods.

When compared with template-free approaches, results shown in Tables 1 and 2 demonstrate that the G2Gs achieves

Table 3. Accuracy of the reaction center prediction module.

Setting	Top- $k$ accuracy %			
	1	2	3	5
reaction class known	90.2	94.5	94.9	95.0
reaction class unknown	75.8	83.9	85.3	85.6

Table 4. Accuracy of the variational graph translation module.

Setting	Top- $k$ accuracy %			
	1	3	5	10
reaction class known	66.8	87.2	91.5	93.9
reaction class unknown	61.1	81.5	86.7	90.0

competitive results on all the cases with different  $k$ . In particular, our G2Gs always outperforms the template-free baselines by a large margin, with up to 63% relative improvement in terms of the top-1 exact match accuracy when the reaction class is known (the second column in Table 1), and up to 29% relative improvement when the reaction class is unknown (the second column in Table 2).

When consider the comparison with template-based baselines, the results in Tables 1 and 2 indicate that our G2Gs approaches or outperforms the state-of-the-art method GLN (Dai et al., 2019), especially when the  $k$  is small. For example, when reaction class is given as a prior, our G2Gs outperforms the GLN in terms of the top-3 and top-5 exact match accuracy.

## 4.3. Ablation Study

To gain insights into the working behaviours of G2Gs, we conduct ablation studies in this section.

**Effectiveness of the Reaction Center Identification Module** We evaluate the top- $k$  accuracy of the reaction center identification module by comparing the true reaction center against the top- $k$  atom pairs with the highest reactivity scores above a threshold. The results in Table 3 indicate that when the reaction class is given as a prior, G2Gs can precisely pinpoint the reaction center in most cases even for  $k = 1$ . When the reaction class is unknown, the performance is slightly worse than the previous case, as a target molecule can usually be synthesized in different routes and thus the model tends to make predictions with low certainty.

## Impact of the Variational Graph Translation Module

To examine the graph translation module, we first split synthons from products based on the true reaction centers (*i.e.*, label matrix  $Y$  in Section 3.2), and then use the same strategy as in Section 4.2 to compute the top- $k$  exact match accuracies. As shown in Table 4, the graph translation mod-

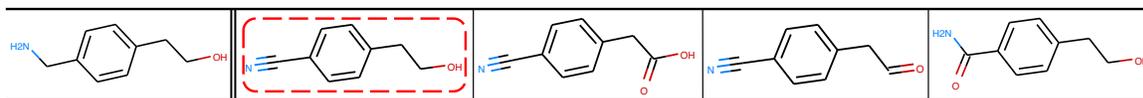


Figure 3. Visualization of the top-1 translation results (right four molecules) of a given product molecule (leftmost), which are conditioned on different latent vectors sampled from prior  $\mathcal{N}(z|0, I)$ . The correct prediction provided is highlighted in red dashed box.

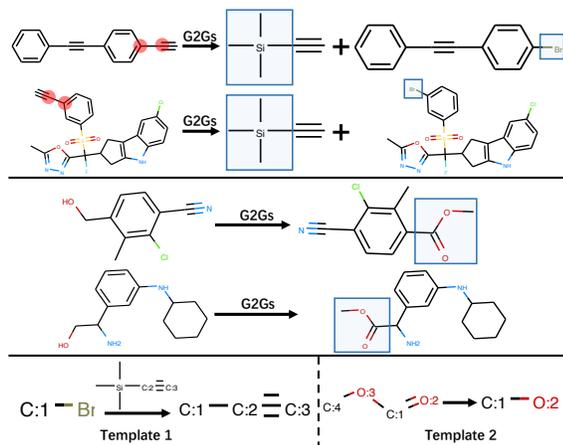


Figure 4. Visualization of several successful cases. Templates for these reactions are summarized at the bottom of the figure.

ule achieves high accuracy on translating the given synthon graphs to reactant graphs, which is an important attribute to the superior performance of the proposed G2Gs framework.

**Diverse Reactant Graph Generation** To observe the benefits brought by the latent vector used in G2Gs, Figure 3 visualizes the top-1 translation results for a given product molecule, which is randomly selected from the test set. In the figure, the right four molecules are generated, based on the same synthon (leftmost molecule in the figure), through randomly sampling from the prior distribution of the formed latent vector. Results in this figure suggest that, through sampling the latent vector formed, our method can generate diverse molecule structures. That is, sampling the learned latent vector, an intrinsic process in G2Gs, powers our model to create diverse and valid reactant graphs, which represents an appealing feature for retrosynthesis prediction.

#### 4.4. Case Study via Visualization

In Figure 4, we show cases where G2Gs successfully identifies the reaction centers and translates the product graph to a set of reactant graphs that match the ground truth. The synthesis routes shown in Figure 4 can be divided into two groups, each of which corresponds to a reaction template presented at the bottom of the figure. These figures suggest that our model does learn domain knowledge from the data set. Such property of our method makes it an appealing solution to practical problems with limited template knowledge.

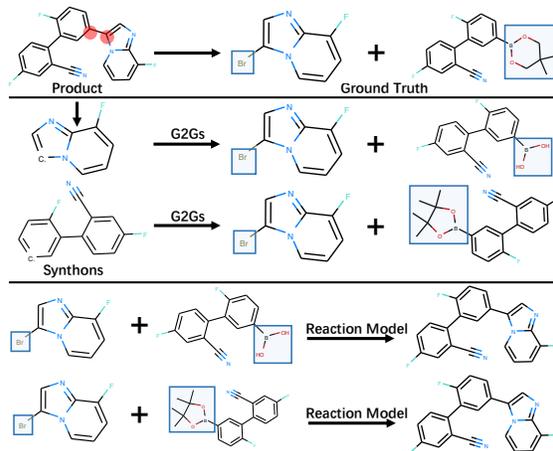


Figure 5. Visualization of a mismatched case. Outcomes of the reactions below the middle line are predicted by the forward reaction prediction model from (Jin et al., 2017).

In Figure 5, we also present a case where none of the predictions matches the ground truth. However, we note that this does not necessarily mean that our model fails to predict a synthesis route for the target molecule. This is because a molecule can be synthesized in multiple ways and the ground truth in the data set is not the only answer. To verify this hypothesis, we employ a forward reaction prediction model (Jin et al., 2017) to predict the product molecules based on the reactants generated by our model. As shown at the bottom of Figure 5, the predicted product exactly matches the target molecule of the retrosynthesis problem. This confirms that the predictions made by G2Gs are indeed potentially valid.

## 5. Conclusion and Outlook

We novelly formulated retrosynthesis prediction as a graph-to-graphs translation task and proposed a template-free approach to attain the prediction goal. In addition, we devised a variational graph translation module to capture the uncertainty and to encourage diversity in the graph translation process. We also empirically verified the superior performance of our proposed method; our strategy outperformed the state-of-the-art template-free counterpart by up to 63% and approached the performance obtained by the state-of-the-art template-based strategy, in terms of top-1 accuracy. Our method excludes domain knowledge and scales well to large datasets, making it particularly attractive in practice.

Our future work will include extending our G2Gs approach to embrace an end-to-end training paradigm and leveraging it to cope with multi-step retrosynthesis tasks (Corey, 1991).

## Acknowledgements

We thank the anonymous reviewers for their instructive feedback. This project is supported by the Natural Sciences and Engineering Research Council of Canada, the Canada CIFAR AI Chair Program, collaboration grants between Microsoft Research and Mila, Amazon Faculty Research Award, Tencent AI Lab Rhino-Bird Gift Fund and a NRC Collaborative R&D Project (AI4D-CORE-06).

## References

- Bahdanau, D., Cho, K., and Bengio, Y. Neural machine translation by jointly learning to align and translate. *arXiv e-prints*, abs/1409.0473, September 2014.
- Baylon, J., Cilfone, N., Gulcher, J., and Chittenden, T. Enhancing retrosynthetic reaction prediction with deep learning using multiscale reaction classification. *Journal of Chemical Information and Modeling*, 59, 01 2019. doi: 10.1021/acs.jcim.8b00801.
- Bengio, S., Vinyals, O., Jaitly, N., and Shazeer, N. Scheduled sampling for sequence prediction with recurrent neural networks. In *Advances in Neural Information Processing Systems*, pp. 1171–1179, 2015.
- Coley, C., Barzilay, R., Jaakkola, T., Green, W., and Jensen, K. Prediction of organic reaction outcomes using machine learning. *ACS Central Science*, 3, 04 2017a. doi: 10.1021/acscentsci.7b00064.
- Coley, C., Rogers, L., Green, W., and Jensen, K. Computer-assisted retrosynthesis based on molecular similarity. *ACS Central Science*, 3, 11 2017b. doi: 10.1021/acscentsci.7b00355.
- Coley, C., Green, W., and Jensen, K. Machine learning in computer-aided synthesis planning. *Accounts of Chemical Research*, 51, 05 2018. doi: 10.1021/acs.accounts.8b00087.
- Corey, E. and Wipke, W. Computer-assisted design of complex organic syntheses. *Science (New York, N.Y.)*, 166: 178–92, 11 1969. doi: 10.1126/science.166.3902.178.
- Corey, E. J. The logic of chemical synthesis: multistep synthesis of complex carbogenic molecules (nobel lecture). *Angewandte Chemie International Edition in English*, 30 (5):455–465, 1991.
- Dai, H., Li, C., Coley, C., Dai, B., and Song, L. Retrosynthesis prediction with conditional graph logic network. In *Advances in Neural Information Processing Systems* 32, pp. 8870–8880. Curran Associates, Inc., 2019.
- Duvenaud, D. K., Maclaurin, D., Iparraguirre, J., Bombarell, R., Hirzel, T., Aspuru-Guzik, A., and Adams, R. P. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in neural information processing systems*, pp. 2224–2232, 2015.
- Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1263–1272. JMLR.org, 2017.
- Hamilton, W., Ying, Z., and Leskovec, J. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, pp. 1024–1034, 2017.
- Hartenfeller, M., Eberle, M., Meier, P., Nieto-Oberhuber, C., Altmann, K.-H., Schneider, G., Jacoby, E., and Renner, S. A collection of robust organic synthesis reactions for in silico molecule design. *Journal of chemical information and modeling*, 51:3093–8, 11 2011. doi: 10.1021/ci200379p.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9:1735–80, 12 1997. doi: 10.1162/neco.1997.9.8.1735.
- Jin, W., Coley, C., Barzilay, R., and Jaakkola, T. Predicting organic reaction outcomes with weisfeiler-lehman network. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems* 30, pp. 2607–2616. Curran Associates, Inc., 2017.
- Karpov, P., Godin, G., and Tetko, I. A transformer model for retrosynthesis. 2019. doi: 10.26434/chemrxiv.8058464.v1.
- Kearnes, S., McCloskey, K., Berndl, M., Pande, V., and Riley, P. Molecular graph convolutions: moving beyond fingerprints. *Journal of computer-aided molecular design*, 30(8):595–608, 2016.
- Kingma, D. and Welling, M. Auto-encoding variational bayes. 12 2014.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations*, 2014.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- Landrum, G. Rdkit: Open-source cheminformatics software. 2016.

- Li, Y., Vinyals, O., Dyer, C., Pascanu, R., and Battaglia, P. Learning deep generative models of graphs. *arXiv preprint arXiv:1803.03324*, 2018.
- Liu, B., Ramsundar, B., Kawthekar, P., Shi, J., Gomes, J., Nguyen, Q., Ho, S., Sloane, J., Wender, P., and Pande, V. Retrosynthetic reaction prediction using neural sequence-to-sequence models. *ACS Central Science*, 3, 06 2017. doi: 10.1021/acscentsci.7b00303.
- Liu, Q., Allamanis, M., Brockschmidt, M., and Gaunt, A. Constrained graph variational autoencoders for molecule design. In *Advances in Neural Information Processing Systems*, pp. 7795–7804, 2018.
- Lowe, D. *Extraction of chemical structures and reactions from the literature*. PhD thesis, 10 2012.
- Neil, D., Segler, M. H. S., Guasch, L., Ahmed, M., Plumbley, D., Sellwood, M., and Brown, N. Exploring deep recurrent models with reinforcement learning for molecule design. In *6th International Conference on Learning Representations, Workshop Track Proceedings*, 2018.
- Olivecrona, M., Blaschke, T., Engkvist, O., and Chen, H. Molecular de-novo design through deep reinforcement learning. *Journal of cheminformatics*, 9(1):48, 2017.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. Automatic differentiation in pytorch. In *NIPS-W*, 2017.
- Popova, M., Shvets, M., Oliva, J., and Isayev, O. Molecular-rnn: Generating realistic molecular graphs with optimized properties. *arXiv preprint arXiv:1905.13372*, 2019.
- Samanta, B., De, A., Ganguly, N., and Gomez-Rodriguez, M. Designing random graph models using variational autoencoders with applications to chemical design. *arXiv preprint arXiv:1802.05283*, 2018.
- Schlichtkrull, M., Kipf, T. N., Bloem, P., Van Den Berg, R., Titov, I., and Welling, M. Modeling relational data with graph convolutional networks. In *European Semantic Web Conference*, pp. 593–607. Springer, 2018.
- Schütt, K. T., Arbabzadah, F., Chmiela, S., Müller, K. R., and Tkatchenko, A. Quantum-chemical insights from deep tensor neural networks. *Nature communications*, 8: 13890, 2017.
- Schwaller, P., Gaudin, T., Lanyi, D., Bekas, C., and Laino, T. "found in translation": Predicting outcomes of complex organic chemistry reactions using neural sequence-to-sequence models. *Chemical Science*, 9, 11 2017. doi: 10.1039/C8SC02339E.
- Schwaller, P., Laino, T., Gaudin, T., Bolgar, P., Bekas, C., and Lee, A. A. Molecular transformer for chemical reaction prediction and uncertainty estimation. *ArXiv*, abs/1811.02633, 2018.
- Segler, M. and Waller, M. Neural-symbolic machine learning for retrosynthesis and reaction prediction. *Chemistry (Weinheim an der Bergstrasse, Germany)*, 23, 01 2017. doi: 10.1002/chem.201605499.
- Segler, M. H., Kogej, T., Tyrchan, C., and Waller, M. P. Generating focused molecule libraries for drug discovery with recurrent neural networks. *ACS central science*, 4 (1):120–131, 2017.
- Shi, C., Xu, M., Zhu, Z., Zhang, W., Zhang, M., and Tang, J. Graph{af}: a flow-based autoregressive model for molecular graph generation. In *International Conference on Learning Representations*, 2020.
- Szymkuć, S., Gajewska, E., Klucznik, T., Molga, K., Dittwald, P., Startek, M., Bajczyk, M., and Grzybowski, B. Computer-assisted synthetic planning: The end of the beginning. *Angewandte Chemie (International ed. in English)*, 55, 04 2016. doi: 10.1002/anie.201506101.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. Attention is all you need. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30*, pp. 5998–6008. Curran Associates, Inc., 2017.
- Weininger, 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*, pp. 6410–6421, 2018a.
- You, J., Ying, R., Ren, X., Hamilton, W. L., and Leskovec, J. Graphrnn: Generating realistic graphs with deep autoregressive models. *arXiv preprint arXiv:1802.08773*, 2018b.