

MIT Open Access Articles

*Polygrammar: Grammar for Digital
Polymer Representation and Generation*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Guo, Minghao, Shou, Wan, Makatura, Liane, Erps, Timothy, Foshey, Michael et al. 2022. "Polygrammar: Grammar for Digital Polymer Representation and Generation." *Advanced Science*.

As Published: 10.1002/advs.202101864

Publisher: Wiley

Persistent URL: <https://hdl.handle.net/1721.1/143799>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Creative Commons Attribution 4.0 International license



Polygrammar: Grammar for Digital Polymer Representation and Generation

Minghao Guo, Wan Shou, Liane Makatura, Timothy Erps, Michael Foshey, and Wojciech Matusik*

Polymers are widely studied materials with diverse properties and applications determined by molecular structures. It is essential to represent these structures clearly and explore the full space of achievable chemical designs. However, existing approaches cannot offer comprehensive design models for polymers because of their inherent scale and structural complexity. Here, a parametric, context-sensitive grammar designed specifically for polymers (PolyGrammar) is proposed. Using the symbolic hypergraph representation and 14 simple production rules, PolyGrammar can represent and generate all valid polyurethane structures. An algorithm is presented to translate any polyurethane structure from the popular Simplified Molecular-Input Line-entry System (SMILES) string format into the PolyGrammar representation. The representative power of PolyGrammar is tested by translating a dataset of over 600 polyurethane samples collected from the literature. Furthermore, it is shown that PolyGrammar can be easily extended to other copolymers and homopolymers. By offering a complete, explicit representation scheme and an explainable generative model with validity guarantees, PolyGrammar takes an essential step toward a more comprehensive and practical system for polymer discovery and exploration. As the first bridge between formal languages and chemistry, PolyGrammar also serves as a critical blueprint to inform the design of similar grammars for other chemistries, including organic and inorganic molecules.

deepen the fundamental understanding, it is extremely beneficial to characterize, enumerate, and explore the entire space of achievable polymer structures. Such a large (ideally exhaustive) collection of polymers would be particularly powerful in conjunction with machine learning and numerical simulation techniques, as a way to facilitate complicated tasks like human-guided molecular exploration,^[1–12] property prediction,^[13–17,69,70,71,72] and retrosynthesis.^[18,19] Computational approaches based on chemical representations and generated data^[20–27] have also tremendously reduced the time, cost, and resources spent on physical synthesis in the chemistry lab.^[28–31]

Ideally, a chemical design model would include three components: 1) a well-defined representation capable of capturing known structures, 2) a generative model capable of enumerating all structures in a given class, and 3) an inverse modeling procedure capable of translating known molecular structures into the representation. For a given class of molecules, an ideal chemical design model should satisfy the

following five criteria: i) Complete: representation is able to encode all possible structures in the given class. ii) Explicit: representation directly specifies the molecular structure. iii) Valid: every generated output is a physically valid chemical structure in the given class. iv) Explainable: the generation process is understandable to the user. v) Invertible: the inverse procedure can translate molecular structures into the given representation. However, designing a chemical model that meets all these criteria is challenging, especially for structurally complex molecules. Most existing approaches are limited to small, simple chemical structures.^[32–36] Even with this limited scope, the design is labor-intensive: the representation language is typically developed first, then extended for generation and inverse modeling. In particular, there have been many systems for molecular line notations^[32,33] and fragment-level description,^[34,35] which were then used as the basis for generative and inverse schemes.^[5–8]

Yet, a comprehensive chemical design model for large polymers remains elusive due to the polymers' inherent complexity. We present a detailed account for each property, including polymer-specific challenges and the performance of existing methods (see Table 1). Since most methods only focus on one

1. Introduction

Polymers are important materials with diverse structure variations and applications. To facilitate customized applications and

M. Guo, W. Shou, L. Makatura, T. Erps, M. Foshey, W. Matusik
Computer Science and Artificial Intelligence Lab
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
E-mail: wojciech@csail.mit.edu

M. Guo
CUHK Multimedia Lab
The Chinese University of Hong Kong
Sha Tin, Hong Kong

 The ORCID identification number(s) for the author(s) of this article can be found under <https://doi.org/10.1002/advs.202101864>

© 2022 The Authors. Advanced Science published by Wiley-VCH GmbH. This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

DOI: 10.1002/advs.202101864

Table 1. Comparison with related chemical design models. Our PolyGrammar is the only approach that satisfies all five properties, simultaneously covering both representation and generative modeling.

Design Models		Representation		Generative Modeling		Inverse Modeling
		Complete	Explicit	Valid	Explainable	Translation from SMILES
Chemical Language	SMILES ^[28]	✓	✓	–	–	✓
	BigSMILES ^[33]	✓	×	–	–	✓
	SELFIES ^[74]	✓	×	–	–	✓
Generation	Bayesian Framework ^[76]	(based on SMILES)		×	×	✓
	GAN ^[9–12]			×	×	✓
	Auto-Encoders ^[5–8,36]			×	×	✓
	STONED ^[75]	(based on SELFIES)		✓	✓	✓
Ours	PolyGrammar	✓	✓	✓	✓	✓

aspect of the chemical design model, i.e. either representation or generative modeling, we separate existing methods into two categories: chemical language and generation. For chemical language, we compare commonly used SMILES^[28] and BigSMILES^[33] together with a newly invented string-based representation SELFIES.^[74] For generation, our comparison covers popular machine learning-based algorithms including Bayesian framework,^[76] generative adversarial network (GAN),^[9–12] and auto-encoders (AE).^[5–8,36] We also list here the STONED algorithm^[75] which generates molecules based on interpolation of SELFIES representation. After exploring the state of the art for all five properties, we give an overview of our proposed approach.

1.1. Complete

Polymers are intrinsically stochastic molecules constructed from some distribution of chemical sub-units. Thus, given a particular set of reactants, the synthesized polymers are not unique; instead, there is wide variation in the resulting structures. For example, consider the polyurethanes synthesized by a 1:1 ratio of two distinct components: methylene bis(phenyl isocyanate) (MDI) and poly(oxytetramethylene) diol (PTMO). Consider one possible outcome of chain length 6, where chain length is defined as the sum of MDI and PTMO units. Disregarding more nuanced chemical restrictions (which are beyond the scope of this paper), any arrangement of the 3 MDI and 3 PTMO units is equally valid. Thus, for a chain of length 20, the component permutations can result in more than $\binom{20}{10} \approx 10^5$ possible structures. This vast set of structures makes it challenging to design a complete and concise polymer representation. Some existing line notations,^[28–34,74] including SMILES^[28] (designed for general molecules), BigSMILES^[33] (specifically designed for large polymers), and SELFIES^[74] are complete representations since they can convert any given polymer structure instance into the form of strings. However, schemes relying on machine learning are not guaranteed to satisfy this property since the learned representation spaces (numeric vectors called latent variables) may exclude polymer structures that do not exist in training data.

1.2. Explicit

The properties of polymeric material are largely determined by the structure of the polymer itself, including the identity and arrangement of its constituent monomers.^[37–39] Thus, it is useful to have an explicit representation for polymers, in which specific structural information is directly expressed and easily understood. This is challenging because a polymer must be understood on many scales, including the overarching structure of repeated units, and the individual molecular and atomic sub-units that comprise them.

Low-level representations like SMILES are able to depict explicit polymeric structures, but the strings are typically hard to parse due to their length. For example, the canonical SMILES representation for the polyurethane chain of length 30 (5 repetitions of the 6-length chain described above) requires more than 600 characters. By contrast, most representations designed for large polymers^[32–34] are so high-level that they are unable to provide explicit information about the complete polymer structure. For example, BigSMILES can express the constituent monomers and the bonding descriptions between them, but it cannot specify the detailed arrangement of the polymer's components. SELFIES uses a sequence of derivations to generate SMILES strings and can only be fully understandable after the string is finalized. As for the machine learning algorithms, the latent variable is an implicit representation and it is impractical to understand the polymer structures merely from the numeric vector.

1.3. Valid

Generative models that build on a well-defined representation scheme are highly coveted,^[40] particularly for their ability to efficiently build large corpora of example structures. However, the result is only useful if the examples generated by the model are guaranteed to be chemically valid. This is challenging to enforce for polymers, as there are many hard chemical constraints (e.g., valency conditions) and other restrictions to account for. The likelihood of violating these constraints increases as the target molecules get larger.

STONED^[75] operates on the constrained SELFIES string space and ensure the generated molecules are valid. On the contrary,

machine learning techniques including support vector machines (SVM),^[41] recurrent neural networks (RNN),^[1–4] generative adversarial networks (GAN),^[9–12] and AE have been used as generative models for molecules. However, these methods often produce chemically invalid outputs, even when limited to small molecules. It is even more challenging for these methods to generate valid polymers, due to the large number of generation steps required to realize such large molecules. Although several recent efforts based on AE^[35,36] and reinforcement learning (RL)^[42,43] have been proposed to produce valid polymers, it is not clear how well they generalize –, i.e., the AE may be unable to ensure validity when generating polymers that significantly deviate from the training data.

1.4. Explainable

To ensure confidence in the results of the generative model, the generation process itself must be fully transparent and understandable to chemists. This property is not necessarily more challenging for large polymers (compared to small molecules), but it is much more critical to facilitate understanding of the resulting polymer structure. Interpretable generation processes also aid the exploration of possible polymer variations.

AE and other deep learning-based generative models^[1–4,9,10,45] produce structures based on implicit latent variables. These models are black-box functions that cannot be easily interpreted. By contrast, the generative model of STONED can be interpreted since its interpolation happens between known molecule structures.

1.5. Invertible

When designing a new chemical design model, it is critical to ensure compatibility with existing notations. In particular, it should be possible (via an inverse modeling procedure) to translate any final representation from an existing scheme into the proposed representation. This inverse procedure should yield the same process and final representation as if the structure were created via the integrated generative model. This is critical for two reasons: i) it makes existing knowledge accessible in the new representation, and ii) it confirms the representative power of the new chemical design model.

To judge invertibility for polymer models, we consider translation from one of the most popular molecule notations: SMILES. As shown in **Table 1**, invertibility is already an important feature common to many existing methods. For example, the encoder of a chemical AE takes a SMILES string as input, then outputs the corresponding latent variable. BigSMILES is built directly upon SMILES so it can easily covert SMILES strings of polymers into the BigSMILES representation. When building our own representation, we also consider “invertibility” with respect to the SMILES format. However, in principle, it is possible to design inverse procedures that translate from other existing representations schemes as well.

1.6. Our Approach

In this paper, we propose a new chemical design model for polymers that respects all five of the ideal properties dis-

cussed above. We introduce PolyGrammar, a parametric context-sensitive grammar for polymers. In formal language theory, grammar describes how to build strings from a language’s alphabet following a set of production rules. PolyGrammar represents the chain structure as a hypergraph. In particular, each polymer chain is represented as a string of symbols, each of which refers to a particular molecular fragment of the original chain. This symbolic hypergraph representation supports explicit descriptions for an infinite amount of diversely structured polymer chains by changing the form of symbolic strings.

Based on this representation, we establish a set of production rules that can effectively generate chemically valid symbolic strings. The recursive nature of grammar production makes it possible to generate any polymer in our given class using only a simple set of production rules. In particular, it is possible for PolyGrammar to enumerate all valid polymers structures within a given class.

As a demonstrative example, we focus on a particular class of polymers: polyurethanes. We choose polyurethanes due to their wide-ranging applications, including antistatic coating,^[46] foams,^[47] elastomers,^[48] and drug delivery for cancer therapy.^[49] Consider generating a polyurethane of chain length of 20, using 1 polyol type (e.g., PTMO) and 1 isocyanate type (e.g., MDI). Under these assumptions (which are representative of the average polyurethane chain^[50]), PolyGrammar can generate more than 2×10^6 distinct polyurethane chains using only 14 production rules. Moreover, we show that PolyGrammar can be easily extended to the other types of polymers, including both copolymers and homopolymers. We further propose an inverse modeling algorithm that translates a polymer’s SMILES string into the sequence of production rules used to generate it. More than 600 polyurethanes collected from literature are validated by this inverse model, demonstrating the representative power of PolyGrammar. The schematic of our PolyGrammar is shown in **Figure 1**.

2. Hypergraph-Based Symbolic Representation

In this section, we introduce the hypergraph representation of polyurethane structures and describe how to use symbolic strings to represent polyurethane chains.

2.1. Polymers as Hypergraphs

It is a common practice^[7,12,51,52] to regard the structural formula of a molecule as an ordinary graph, where atoms are nodes, bonds are edges, and edges connect exactly two nodes. For polyurethanes, ordinary graph depictions would require prohibitively many nodes and edges. To address this, we employ a generalized graph called a hypergraph,^[53] which allows individual edges to join more than one node. Any edge that connects a subset of the nodes in the hypergraph is called a hyperedge. Consider the product of two monomers (1,3 bis(isocyanatomethyl)cyclohexane and diethylene glycol) as shown in **Figure 2i**. Originally, the graph requires 21 nodes and 21 edges. However, if we construct each hyperedge by selecting the subset of nodes according to the monomer type, as shown in

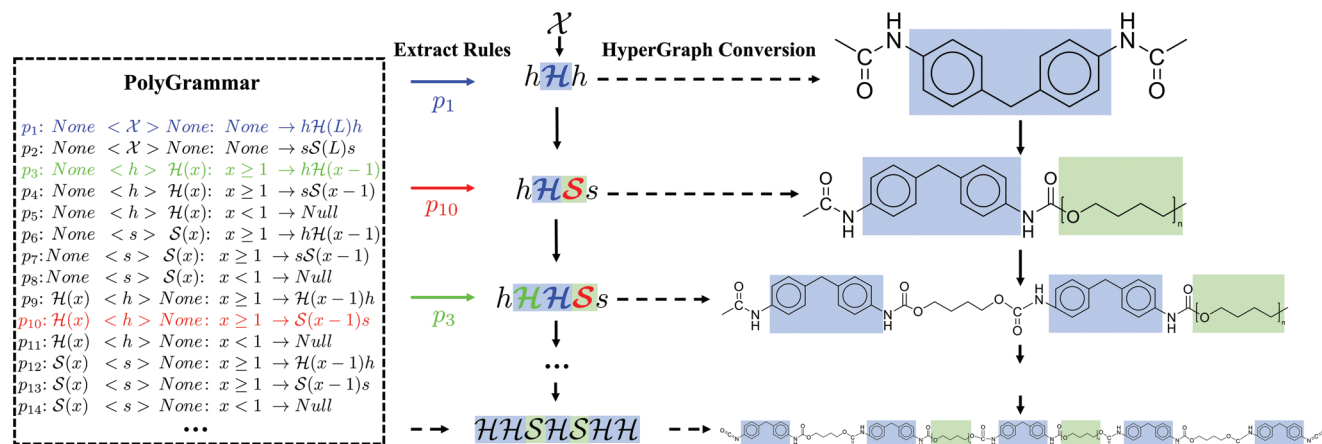


Figure 1. Schematic of our chemistry design model, PolyGrammar, which represents molecular chain structure as a string of symbols (center). PolyGrammar consists of a set of production rules $\{p_i | i = 1, \dots, 14\}$ (left). The generation process starts from an initial symbol \mathcal{X} . At each iteration, each non-terminal symbol (h , s or \mathcal{X}) in the current string is replaced by the successor of a production rule whose predecessor matches the symbol. The generation process concludes when the string does not contain any non-terminal symbols. The resulting symbol string (center) is then translated to a polymer chain (right) by hypergraph conversion.

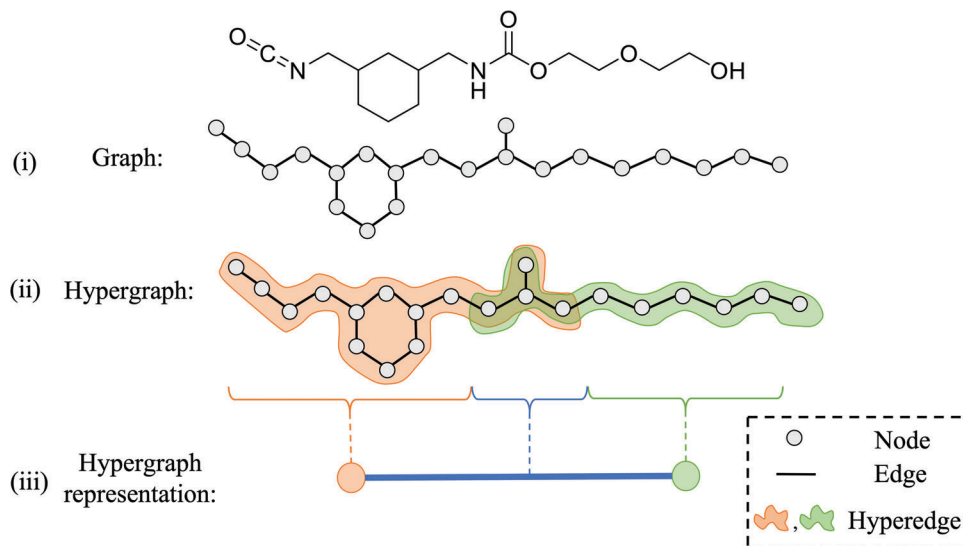


Figure 2. The structure produced by reacted by two monomers (1,3-bis(isocyanatomethyl)cyclohexane and diethylene glycol). The standard graph representation i) uses 21 nodes and 21 edges, but the hypergraph ii) only requires two hyperedges. Each hyperedge corresponds to the nodes of a given monomer. Both hyperedges have the urethane group in common. We use line graph iii) to visualize the hypergraph representation in the remaining figures of the paper for convenience.

Figure 2ii, the hypergraph for this molecule requires only 2 hyperedges. This dramatically reduces the representation cost for large polyurethane chains.

For increased convenience, we will visualize the hypergraph representations using the line graph^[54] form shown in Figure 2iii. In graph theory, the line graph refers to the duality of the original graph, where each edge in the original graph corresponds to a unique vertex of the line graph. With regards to the theory of hypergraph, the line graph contains one vertex for every hyperedge in the original hypergraph. Two vertices in the line graph are connected by a line if their corresponding hyperedges in the original hypergraph have a non-empty intersection. For the hypergraph in Figure 2ii, since the urethane group is shared

by two hyperedges in the hypergraph, the corresponding line graph can be visualized as two vertices connected by an edge. By collapsing the original nodes based on molecular identity, the line graph form provides a more concise visualization of a hypergraph.

Complete polyurethane structures can also be represented in this manner. The molecular fragments corresponding to the isocyanate and the polyol in the polyurethane chain are represented as hyperedges, which are visualized as vertices in the line graph. The urethane groups connecting hard segment (HS) with soft segment (SS) and the chain extenders connecting two diisocyanates are viewed as intersections between two hyperedges; thus, they are visualized as edges in the line graph. Two examples

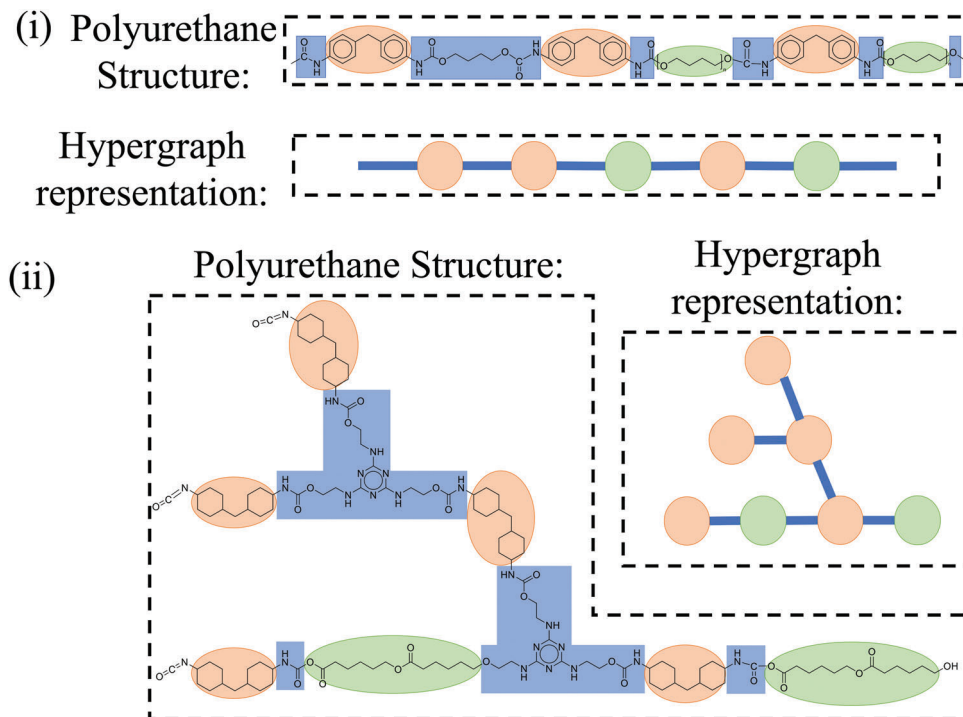


Figure 3. Examples of hypergraph representation. i) Polyurethane chain synthesised by MDI, PTMO, and 1,4-butanediol (BDO); ii) Branched polyurethane chain synthesized by 4,4'-diisocyanato-methylenedicyclohexane (4,4'-HMDI), poly(caprolactone) diol (PCL), and tri-azine based polyhydric alcohol (3-THA).

of hypergraph representations for polyurethane structures are shown in **Figure 3**.

2.2. Symbolic Representation

Given the hypergraph of a polyurethane chain, we construct a corresponding symbolic string for use in PolyGrammar. In the symbolic string, the hyperedges corresponding to the isocyanate (hard segment) are denoted with “ H ” and those corresponding to the polyol (soft segment) are denoted as “ S ”. The chain extenders are omitted since they can only exist between two adjacent H (or S) symbols. For those polyurethanes containing multiple isocyanate or polyol types, we use subscripts $i = 1, 2, \dots$ to distinguish different subtypes of certain hyperedge. For instance, if two different types of isocyanates are used,^[38] we use H_1 and H_2 to distinguish the hyperedges corresponding to each hard-segment type. These rules allow us to represent any polyurethane chain as a string of symbols. Examples are shown in **Figure 4**.

We emphasize that our symbolic representation is invertible, such that a symbolic string can be converted back to the corresponding chemical structure if the constituent isocyanate(s), polyol(s), and chain extender(s) are specified. We call this process hypergraph conversion. The invertibility of hypergraph representation ensures our PolyGrammar can simultaneously serve as a representation and a generative model for polyurethanes.

3. PolyGrammar

In this section, we first present the basic mechanism of grammar production using an illustrative example. Then, we introduce our

parametric context-sensitive PolyGrammar comprehensively. Finally, we propose several advanced features based on our basic PolyGrammar for the representation of polyurethanes, which encourage the generation of more general structures.

3.1. Basic PolyGrammar

In formal language theory, a grammar $G = (N, \Sigma, P)$ is used to describe a language, where N is a set of non-terminal symbols, Σ is a set of terminal symbols and P is a set of production rules, each of which consists of a predecessor and a successor separated by a right arrow “ \rightarrow ”. In the language represented by the grammar G , each word is a finite-length string containing both terminal and non-terminal symbols. The non-terminal symbols in a word can be further replaced and expanded by invoking one production rule from P at a step. In our PolyGrammar, the set of non-terminal symbols N is $\{\mathcal{X}, h, s\}$ and the set of terminal symbols Σ is $\{H, S\}$. **Figure 5** shows an illustrative example to demonstrate the process for producing a string via the grammar G . This example uses four production rules: $P = \{\mathcal{X} \rightarrow Hh; h \rightarrow Hh; h \rightarrow Ss; s \rightarrow \text{Null}\}$. Starting from the initial symbol \mathcal{X} , at each iteration, each non-terminal symbol in the current string is replaced with the successor of a production rule whose predecessor matches the symbol. The process continues until no non-terminal symbols exist in the string.

According to Chomsky’s classification,^[55] the grammar used in this illustrative example is a Type-2 grammar, also called context-free grammar, where the predecessor of each production rule consists of only one single non-terminal symbol. Similar

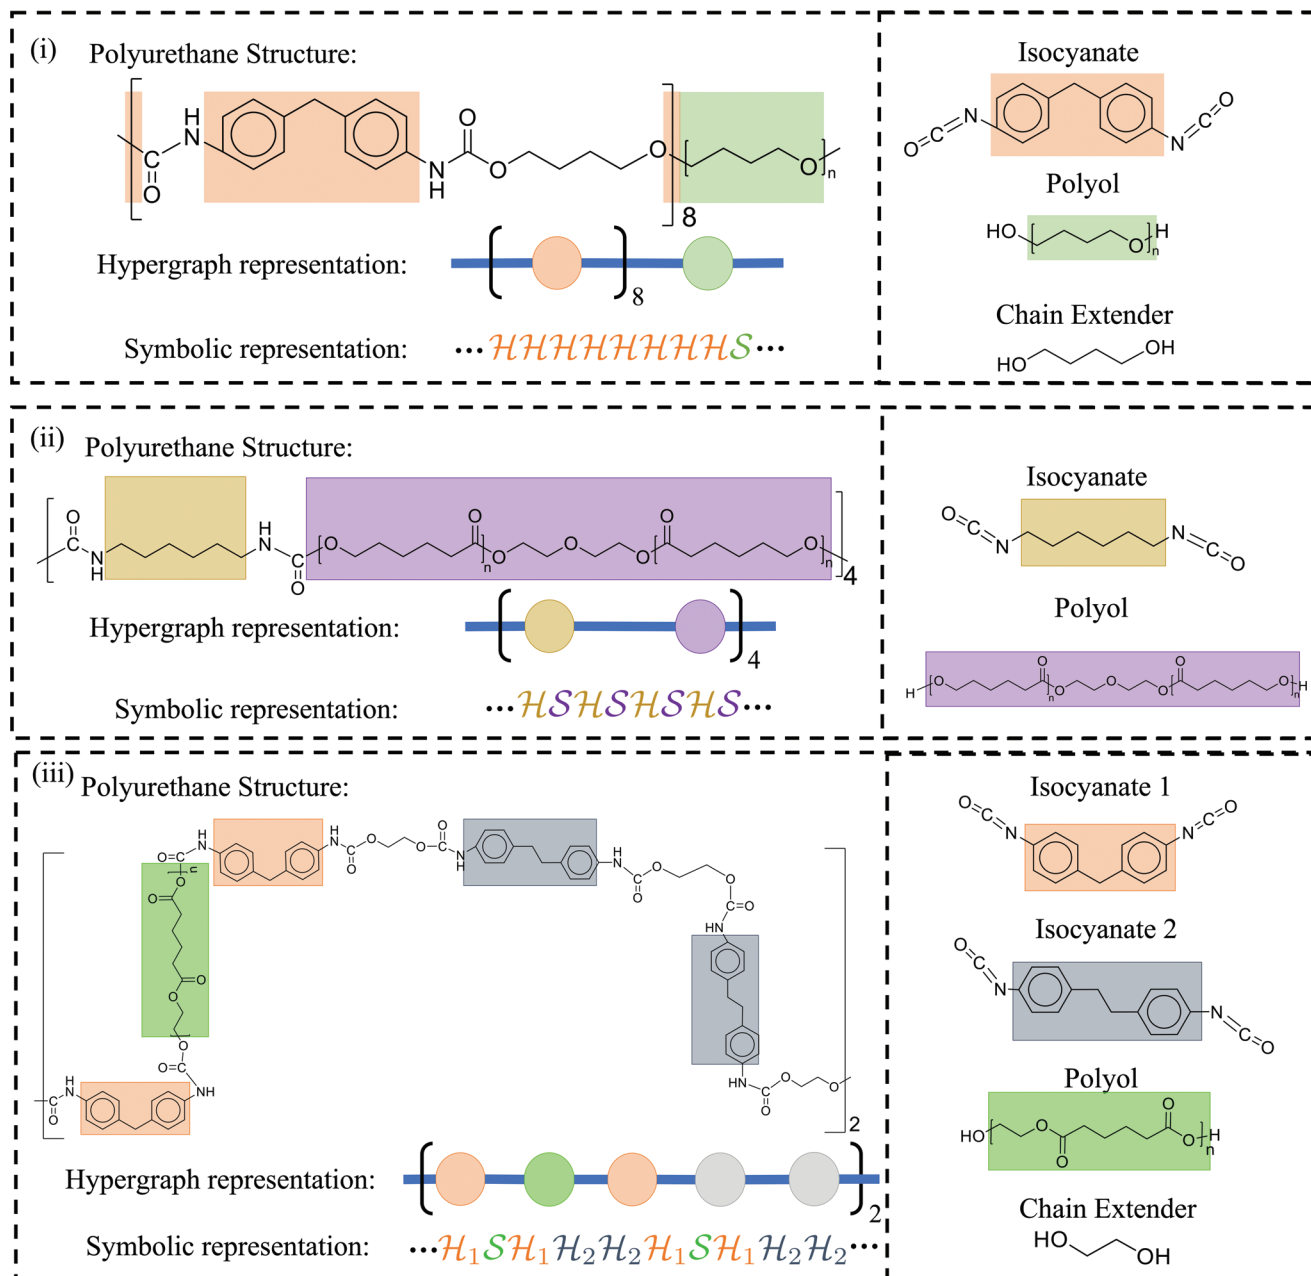


Figure 4. Symbolic representations for polyurethanes synthesized using: i) MDI, PTMO, and BDO; ii) 1,6-diisocyanatohexane (HDI), and PCL; iii) 4,4'-diphenyl diisocyanate (DBDI), MDI, poly(ethylene adipate)diol (PEA), and ethylene glycol (EG). Note that (iii) includes multiple diisocyanates.

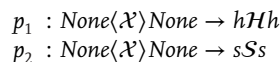
paradigms are also utilized in L-systems to model the morphology of organisms.^[56,57]

3.1.1. Context-Sensitive Grammar

The context-free grammar discussed above is insufficient to imitate the polyurethane generation process because the symbolic string can only expand along one direction; however, polyurethanes generally grow along two opposite directions to form chain structures. To address this, our PolyGrammar uses

context-sensitive grammar. In particular, our PolyGrammar is a Type-1 grammar, a more general form of Type-2 grammar,^[58] where the production rules also consider the context (i.e., the surrounding symbols) of the given non-terminal symbol within the string.

By considering the symbol contexts, the production rules of a context-sensitive grammar can explicitly depict the growing direction of the polyurethane chain. The production rules are as follows:



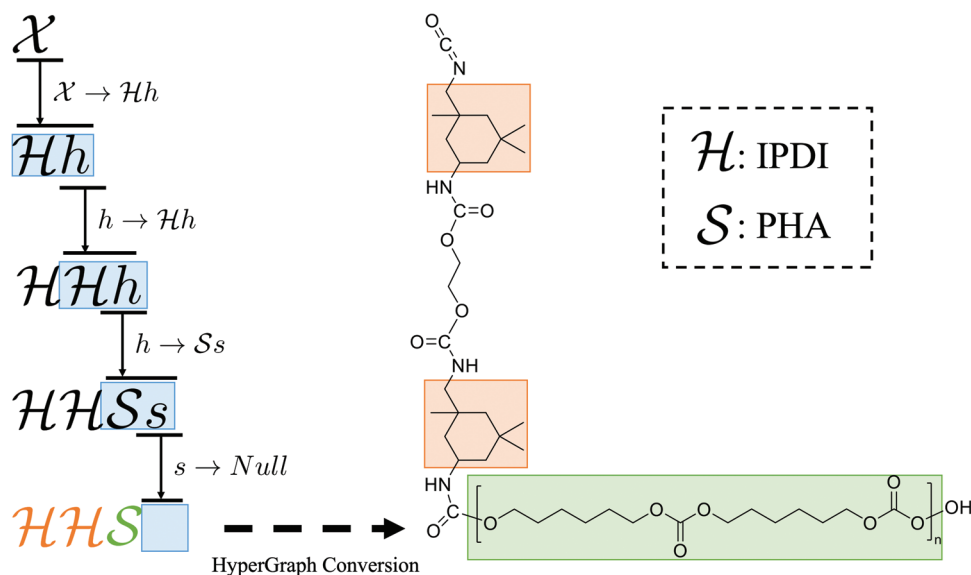


Figure 5. An illustrative example of grammar production. Starting from the initial symbol \mathcal{X} , we sequentially invoke four production rules from $P = \{\mathcal{X} \rightarrow Hh; h \rightarrow Hh; h \rightarrow Ss; s \rightarrow \text{Null}\}$. The process continues until all symbols in the string are terminal symbols. By specifying the constituent structures, i.e., isophorone diisocyanate (IPDI), polyhexamethylene carbonate glycol (PHA), and EG, the string of the symbols can be translated to the corresponding polyurethane chain via hypergraph conversion.

- $p_3 : \text{None} \langle h \rangle H \rightarrow hH$
- $p_4 : \text{None} \langle h \rangle H \rightarrow sS$
- $p_5 : \text{None} \langle h \rangle H \rightarrow \text{Null}$
- $p_6 : \text{None} \langle s \rangle S \rightarrow hH$
- $p_7 : \text{None} \langle s \rangle S \rightarrow sS$
- $p_8 : \text{None} \langle s \rangle S \rightarrow \text{Null}$
- $p_9 : H \langle h \rangle \text{None} \rightarrow Hh$
- $p_{10} : H \langle h \rangle \text{None} \rightarrow Ss$
- $p_{11} : H \langle h \rangle \text{None} \rightarrow \text{Null}$
- $p_{12} : S \langle s \rangle \text{None} \rightarrow Hh$
- $p_{13} : S \langle s \rangle \text{None} \rightarrow Ss$
- $p_{14} : S \langle s \rangle \text{None} \rightarrow \text{Null}$

In each production rule, the non-terminal symbol to be replaced is inside the angle brackets “< >” of the predecessor. The contexts are the symbols located at both sides of “< >” in the predecessor (*None* indicates no constraints). The rule can only be deployed when both contexts of the symbol have been matched.

Each rule has an intuitive function. Rules p_1 and p_2 initialize the start symbol \mathcal{X} , while p_5 , p_8 , p_{11} , and p_{14} terminate the growth. Rules p_3 , p_4 , p_6 , and p_7 extend the string along the left direction, and p_9 , p_{10} , p_{12} and p_{13} extend the string along the right direction. p_3 and p_9 indicate the reaction between two isocyanates, imitating the formation of the hard segment, while p_7 and p_{13} indicate the reaction between two polyols, imitating the formation of the soft segment. Lastly, p_4 , p_6 , p_{10} , and p_{12} imitate the formation of the urethane group.

Another important feature of the PolyGrammar is that there are multiple possible production rules to expand a given symbol. For instance, p_3 , p_4 , and p_5 share the same predecessor and expand the non-terminal symbol h along the left direction. There are many possible schemes for selecting among these options, including hand-tuned heuristics or manual intervention to guide the scheme toward particular results. For simplicity, we have implemented a uniformly random selection technique: at each itera-

tion, we randomly sample one rule from all of the candidate rules that meet the contexts and apply it to the symbol. An example of the production process is illustrated in **Figure 6**.

3.1.2. Parametric Grammar

Although the context-sensitive grammar makes it possible to generate a variety of polyurethane chain structures, its modeling power is still limited. One important problem is that the total chain length of the generated polyurethanes cannot be controlled. In practice, the chain length is an essential factor that influences the physical and chemical properties of polyurethanes.^[50,59] It is non-trivial to control the chain length of each generated polyurethane merely using the grammar discussed above due to the stochastic production. In order to address this problem, we introduce a parameter x associated with each terminal symbol in the grammar and augment our PolyGrammar as a parametric context-sensitive grammar. The proposed parametric grammar is illustrated as follows,

- $p_1 : \text{None} \langle \mathcal{X} \rangle \text{None} : \text{None} \rightarrow hH(L)h$
- $p_2 : \text{None} \langle \mathcal{X} \rangle \text{None} : \text{None} \rightarrow sS(L)s$
- $p_3 : \text{None} \langle h \rangle H(x) : x \geq 1 \rightarrow hH(x-1)$
- $p_4 : \text{None} \langle h \rangle H(x) : x \geq 1 \rightarrow sS(x-1)$
- $p_5 : \text{None} \langle h \rangle H(x) : x < 1 \rightarrow \text{Null}$
- $p_6 : \text{None} \langle s \rangle S(x) : x \geq 1 \rightarrow hH(x-1)$
- $p_7 : \text{None} \langle s \rangle S(x) : x \geq 1 \rightarrow sS(x-1)$
- $p_8 : \text{None} \langle s \rangle S(x) : x < 1 \rightarrow \text{Null}$
- $p_9 : H(x) \langle h \rangle \text{None} : x \geq 1 \rightarrow H(x-1)h$
- $p_{10} : H(x) \langle h \rangle \text{None} : x \geq 1 \rightarrow S(x-1)s$
- $p_{11} : H(x) \langle h \rangle \text{None} : x < 1 \rightarrow \text{Null}$
- $p_{12} : S(x) \langle s \rangle \text{None} : x \geq 1 \rightarrow H(x-1)h$
- $p_{13} : S(x) \langle s \rangle \text{None} : x \geq 1 \rightarrow S(x-1)s$
- $p_{14} : S(x) \langle s \rangle \text{None} : x < 1 \rightarrow \text{Null}$

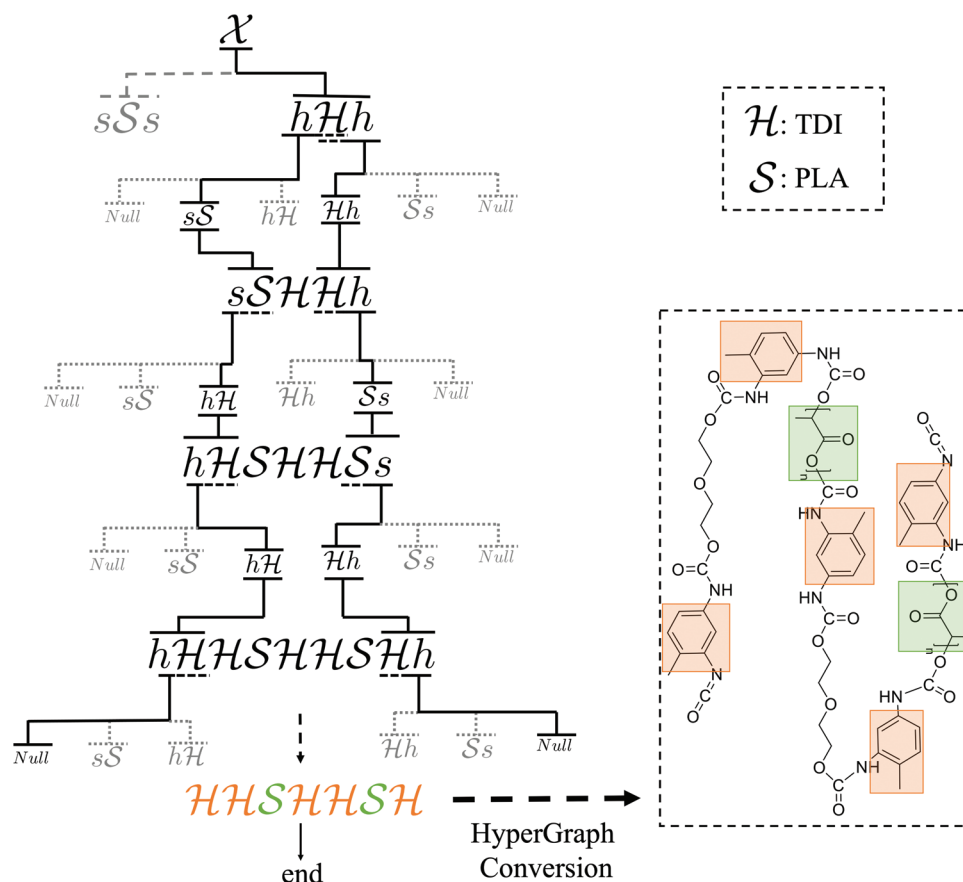


Figure 6. Example of context-sensitive grammar. At each production step, only the rules that match the non-terminal symbol's context are adopted. Hence, the production process can explicitly depict the growing direction of the polyurethane chain. If there are multiple candidate rules at a given step, selection can be done manually or randomly. The selected rule is then applied to the symbol to continue production.

The production rules now feature parameters, which are denoted with parentheses “()” following terminal symbols. Furthermore, each production rule is augmented with a logical “condition” that determines whether the rule can be invoked or not (None indicates no constraints). By specifying L (the initial value of parameter x in production rules p_1 and p_2), the grammar can produce strings with length $2L + 1$, corresponding to polyurethane chains with length $2L + 1$. By varying the value of L , the chain length of generated polyurethanes can be controlled. An example of this production process is illustrated in **Figure 7**.

3.2. Advanced Features

3.2.1. Extensions for Branched Polyurethanes

So far, all of our polyurethanes have featured linear chain structures. However, it is possible for polyurethanes to have branched structures,^[60] as shown in Figure 3ii. To generate branched polyurethanes, we augment the parametric context-sensitive grammar with several rules:

$$p_{15} : \text{None} \langle h \rangle \mathcal{H}(x) : x \geq 1 \rightarrow h[h]\mathcal{H}(x-1)$$

$$p_{16} : \text{None} \langle h \rangle \mathcal{H}(x) : x \geq 1 \rightarrow s[s]\mathcal{S}(x-1)$$

$$p_{17} : \mathcal{H}(x) \langle h \rangle \text{None} : x \geq 1 \rightarrow \mathcal{H}(x-1)[h]h$$

$$p_{18} : \mathcal{S}(x) \langle s \rangle \text{None} : x \geq 1 \rightarrow \mathcal{S}(x-1)[s]s$$

A branch is delimited by the content inside a pair of square brackets “[]”. The non-terminal symbols inside the square brackets can also be further expanded using the rules of the basic PolyGrammar. In the final string, all the terminal symbols inside a pair of square brackets together form a sub-branch attached to the backbone. The above-illustrated rules can generate polyurethane chains that have up to 2 branches at each bifurcation. The number of branches at each bifurcation can also exceed 2 by adding more square-bracket pairs attached to the non-terminal symbols. Examples are available in Supporting Information.

3.2.2. Extensions for Meta-Ring structures

For now, our PolyGrammar focuses on single-chained molecular structures. However, synthesized polyurethanes are a mixture of differently structured chains, where interactions between chains such as hydrogen bonding and crosslinking may occur.^[47,48] These interactions influence the physical and chemical properties of the polyurethane, largely determining whether the synthesized polyurethane is thermoset or thermoplastic. We further

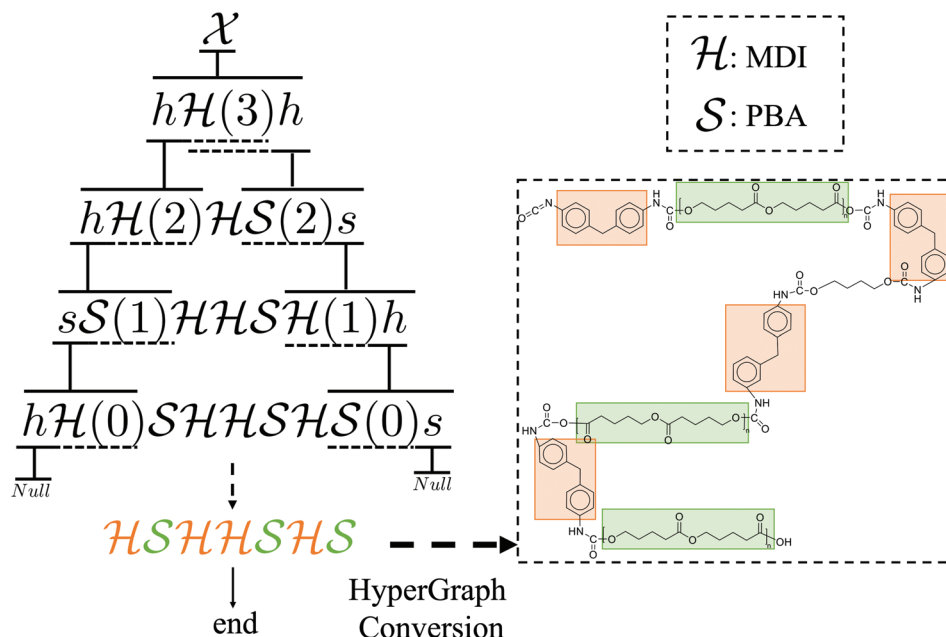


Figure 7. Example of parametric grammar. To control the length of generated polyurethane, we introduce parameters, denoted with parentheses “()” after terminal symbols.

propose a graph grammar^[68] based on the initial PolyGrammar by augmenting the production rules to support interactions between multiple chains. The key idea is to enable a certain production rule to have a simple ring structure at the right-hand side. This ring structure contains non-terminal symbols, which can be further expanded using other production rules to form a larger ring. Since this ring expansion rule can be selected multiple times during the production process, it is possible to have multiple meta-rings in the final generated symbolic graph. By properly arranging the symbols, the graph is isomorphic to multiple chains with interactions between each other. We can then perform hypergraph conversion by specifying the fragment and interaction type to get the final polymer microstructure, including hydrogen bonding and crosslinking. Detailed rules and explanations together with an example of the whole production process for a polymer network^[73] formed by cross-linking poly-(4-vinyl pyridine) (P4VP) with bis-Pd (II) complexes are available in Supporting Information.

3.2.3. Global Controllable Parameters

We have already discussed the use of parameters for controlling the chain length of the generated polyurethanes. However, it is still difficult for our baseline parametric grammar to achieve more advanced controllable parameters such as the ratio of hard segment to soft segment. This is because the context-sensitive grammar only captures “local” information about the chain during the generation process, as the view of each production rule is limited to the context immediately surrounding the predecessor symbol. When it comes to global constraints, such as specific ratios of hard versus soft segments, the generative model

needs to be aware of the relevant information (number of hard segments, chain length) over the whole chain. It is non-trivial to handle these constraints with the basic PolyGrammar discussed in previous sections.

To address this issue, we introduce an additional symbol “*M*” which serves as a message that can collect global information about the chain. The message is propagated back and forth between the left and right ends of the string. The propagation is achieved by switching the message’s position with the adjacent symbol’s one at a time; this continues along a certain direction until the message gets to the string end. At each position swap, the message updates its parameters to collect the information required for the control setting. When the message reaches the end of the string, the outcome of the production rule is influenced by the information contained in the message. The message is then reset and begins to propagate along the opposite direction, encoding information about the entirety of the structure, continuing the above process. Since the production rules are only applied at the end of the chain, this mechanism ensures that the string generation adheres to all parameter-controlled constraints. Multiple constraints can be considered simultaneously by adding more parameters to the message symbol. The full set of the production rules and an illustration of the message passing mechanism are shown in Supporting Information.

4. PolyGrammar as a Generative Model

Generative models are critical for the efficient, thorough exploration of possible polymer structures. These models are also particularly powerful in conjunction with machine learning algorithms, in order to address complicated problems like human-guided exploration and property prediction. In this section, we

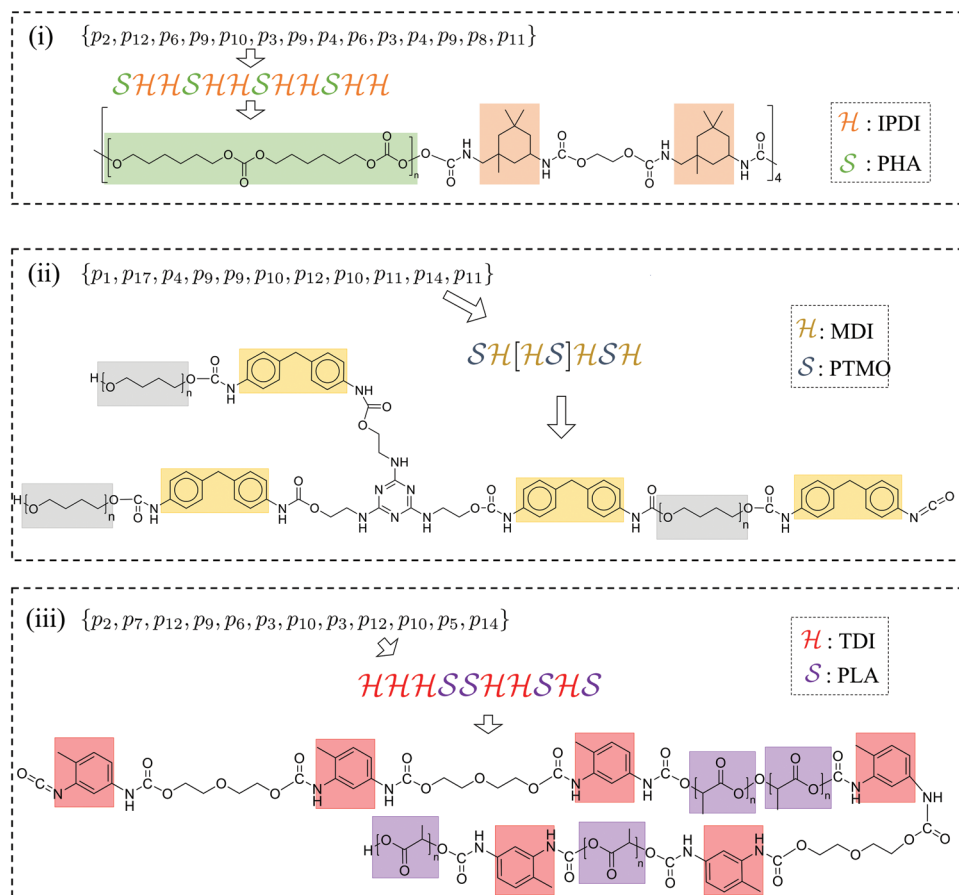


Figure 8. Examples of polyurethane chains generated using PolyGrammar. i) Ordered chain with isophorone diisocyanate (IPDI), polyhexamethylene (PHA) and EG; ii) Branched chain with MDI, PTMO and 3-THA; iii) Unordered chain with Toluene diisocyanate (TDI), PLA and diethylene glycol (DEG).

discuss how our parametric context-sensitive PolyGrammar can serve as a generative model.

The generation process of PolyGrammar begins with a simple string that contains the initial symbol \mathcal{X} . On each step, we traverse the symbols in the current string and find the position of all the non-terminal symbols. For each non-terminal symbol, we identify a candidate set of production rules. Each candidate production rule must meet the following conditions: 1) the context in the predecessor clause matches the context of the current symbol in the string, and 2) the parameters of this symbol's context meet the logical condition of the production rule. If there are several candidate rules to expand a given symbol, a single rule is selected according to the desired scheme (random sampling scheme, manual intervention, etc.). We apply the selected production rule to the appropriate non-terminal symbol, and repeat this process until no non-terminal symbols remain in the string. Once the final string is produced, we convert it into an explicit polyurethane hypergraph by replacing the symbols in the string with the chemical structures (e.g., MDI and PTMO) corresponding to each hyperedge. This yields a valid, explicit polyurethane chain, as desired. These structures can be further converted to other forms of representation such as SMILES.

Using our generative model, it is possible to enumerate all valid polyurethane structures in a target class (e.g., length 20 with

1 type of polyol and 1 type of isocyanate). In particular, any distinct sequence of production rules on the start symbol yields a distinct string, which in turn represents a unique polyurethane chain. Since the production rules encode all permissible local configurations of the constituent molecules, it follows that our grammar is able to generate any valid polyurethane.

To emphasize the volume of achievable molecules, we also quantitatively analyze the diversity of generated chains for our PolyGrammar. Given a chain length parameter L and the number of isocyanate and polyol types (N_H and N_S , respectively), the basic PolyGrammar (with 14 production rules) allows the generation of a total number of

$$N = \sum_{i=0}^{2L+1} \frac{(2L+1)!}{i!(2L+1-i)!} N_H^i N_S^{2L+1-i} \quad (1)$$

polyurethane chains with different structures. With $L = 10$, $N_H = 1$, $N_S = 1$, which are representative of an average polyurethane chain,^[50] N is more than 2×10^6 . This demonstrates the powerful capacity of our PolyGrammar. Several polyurethane chains generated using PolyGrammar are shown in **Figure 8**. More examples can be found in Supporting Information.

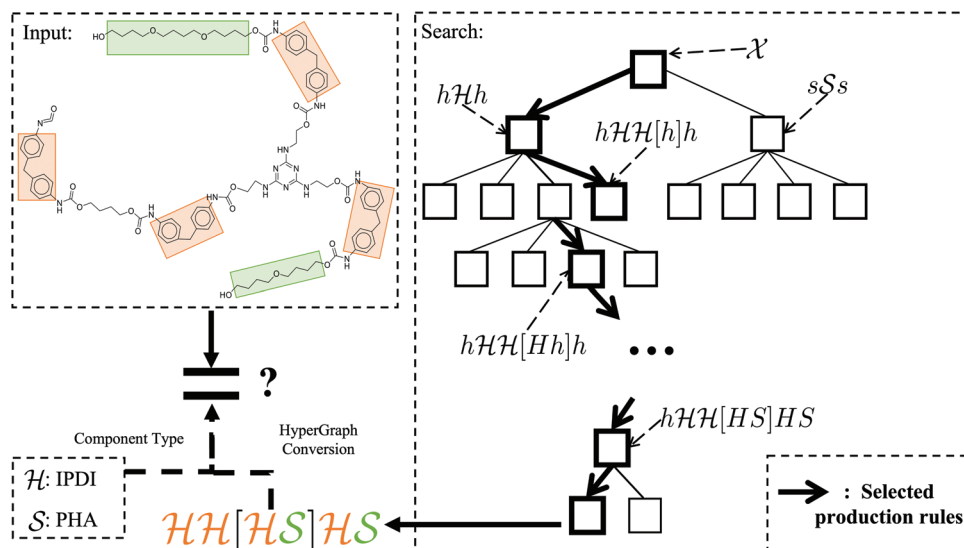


Figure 9. Schematic for translating a polyurethane from a SMILES string into our PolyGrammar representation, which also reveals the complete sequence of rules required for its generation. The pipeline can be regarded as a search process. Starting from the initial symbol, we iteratively select and invoke production rules until all symbols in the string are terminal symbols. Then given the component types, we convert the symbolic string into a polyurethane structure by hypergraph conversion and compare it with the input structure. The total process repeats until the search structure matches the input structure. Note that the component type is not a necessary input of the total algorithm. It can be replaced by another search process in a monomer dataset collected from the literature.

5. Translation from SMILES

To complete our chemical design model, we also develop an inverse model capable of translating a SMILES string into the corresponding sequence of PolyGrammar production rules. The overall pipeline of translation from SMILES can be regarded as a search process, as shown in **Figure 9**. Starting from the initial symbol, we iteratively select and invoke production rules until all symbols in the string are terminal symbols. Once we have a complete string and the specific component types, we use hypergraph conversion to convert the symbolic string into a polyurethane structure. We then compare our results with the input structure; if they do not match, we restart our search from scratch. The process repeats until our structure matches the original input. Note that the component type is not a necessary input of the total algorithm. It can be replaced by another search process in a monomer dataset collected from the literature. This monomer dataset is provided in the Supporting Information.

Specifically, our inverse model proceeds as follows. Given the SMILES string of the polyurethane chain, we break it into multiple molecular fragments by disconnecting all of the urethane groups, $-NHCO-O-$. Then we exhaustively enumerate each molecular fragment and perform a string matching algorithm (KMP matching^[63]) to identify the type of it: an isocyanate, a polyol, or a chain extender. During the enumeration, we also record the connectivity between each fragment. Based on the types and the connectivity of the fragments, we can obtain a hypergraph representation of the original SMILES string. The final step is to convert the hypergraph into the sequence of the production rules of PolyGrammar. We traverse the hypergraph using the breadth-first search (BFS) algorithm, which explores all

of the neighboring hyperedges at the present depth before moving on to the nodes at the next depth level. BFS starts at the tree root, which is an arbitrary hyperedge of the hypergraph. Each step of the exploration returns a tuple of two hyperedges, which is then matched with a specific production rule in the PolyGrammar. Hence, the sequence of the production rules can be obtained once the entire hypergraph has been explored. The pipeline of this algorithm is illustrated in **Figure 10** and the corresponding pseudo-code is in Supporting Information.

This pipeline is sufficient for our needs, but it could be improved with a heuristic search such as A^* search,^[64] best-first search,^[65] or learned heuristic search^[66] where a heuristic function accelerates the search process by directing attention toward the most promising regions of the search space.

To validate our approach and demonstrate the capacity of our proposed PolyGrammar, we have collected and inversely modeled over 600 polyurethane structures from the literature. Many of these polyurethanes are commonly used in synthesis and real-world fabrication, and they feature a wide range of constituent molecules. In particular, the dataset features 8 different types of isocyanates, 11 types of polyols, and 7 types of chain extenders. Additional details about our dataset – including information about how to add and translate new polyurethane structures – are described in Section 7. Supporting Information also contains several examples of polyurethanes from our dataset, which were successfully converted from SMILES to the PolyGrammar representations. Moreover, we emphasize that each of the collected SMILES strings in our dataset can be successfully converted to a sequence of production rules in the PolyGrammar. This proves that our PolyGrammar has a high representative capacity over a large span of polyurethane structures.

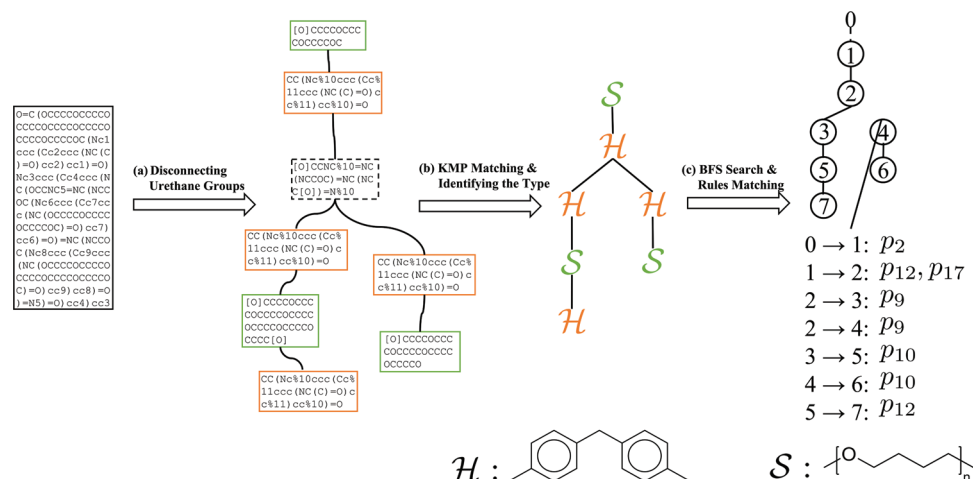


Figure 10. Overview of the algorithm for translation from SMILES. The input SMILES string is first broken into a set of molecular fragments, which are identified via string matching. Based on the identity and connections of each fragment, we can construct the hypergraph representation of the molecule. Then, we search for a sequence of PolyGrammar rules that yields the desired result.

6. Generalization to Other Polymers and Stereochemistry

Our PolyGrammar can also be easily extended to new classes of polymers. These extensions would use the same framework described above, with very few modifications. In the Supporting Information, we illustrate the extended PolyGrammar for different types of copolymers, including alternating copolymers and block copolymers. Note that our PolyGrammar in the main paper can already cover random copolymers, branched copolymers, and graft copolymers. Users only need to add new types of reactants to the symbolic representation in order to determine the species of monomer.

For now, PolyGrammar focuses on the backbone structure, i.e., the arrangement of monomers, which largely determines the property of copolymers (derived from more than one species of monomer). The grammar treats the monomer fragment as a whole and distinguishes different monomer types using different symbols. However, there is also a wide range of polymers consisting of only one single type of repeat unit, i.e., homopolymers, where the backbone structures are not variable and the functional group (also called functional residue) of the monomer contributes to the polymer property. To handle this, we augment our PolyGrammar with an additional set of production rules focusing on the representation and generation of functional groups. We also demonstrate the effectiveness of our augmented PolyGrammar using polyacrylate as an illustrative example. This functional-group grammar together with the basic PolyGrammar (full set of the production rules in Supporting Information) serves as a hierarchical generative model for polymers, where the latter one handles the backbone and the former one focuses on the functional residue of each composed monomer. More examples are shown in Supporting Information.

We also show in Supporting Information that PolyGrammar can handle stereochemistry of polymers by adding an additional parameter “*t*” as the orientation indicator. We use binary numbers “0” and “1” as the parameter to distinguish two different

oriented units. By specifying the logical conditions of each rule, we can control the final generated polymers to hold different tactics. For example, syndiotactic polymers can be obtained if each production rule alters the binary parameter, while atactic polymers can be obtained if each production rule randomly samples the parameter. We show the detailed rules for three common tacticity settings and use polypropylene as an illustrative example in Supporting Information. A similar binary parameter approach can be used to represent charged polymer chains, including polyelectrolytes, where we can use “ $q = 1$ ” for those fragments with positive charges and “ $q = 0$ ” for those with negative charges.

7. Statistical Analysis

We have collected a dataset of polyurethanes from the literature, including 8 different types of isocyanates, 11 types of polyols, and 7 types of chain extenders. Each sample is illustrated in the form of BigSMILES (see Supporting Information for details).

By combining 3 types of components, this dataset contains $8 \times 11 \times 7 = 616$ types of polyurethanes that are commonly used in synthesis and real-world fabrication. The full names of the abbreviations in the dataset are listed in Table S5 (Supporting Information). These data samples are stored in a “.CSV” file and can be easily handled using Python code to perform the algorithms of generative model and translation from SMILES. It is also capable of adding new structures to this dataset. The only thing to do is convert the structure to the BigSMILES format and add it to the “.CSV” file.

8. Discussion

PolyGrammar is an effective chemistry design model that satisfies all five desirable properties discussed in the Introduction. In particular, our symbolic representation can convey all possible polyurethane structures in an explicit yet concise manner. The generative model based on this representation is exhaustive (it

is capable of generating any polyurethane) and trustworthy (every generated polyurethane is guaranteed to be valid). Moreover, the generation process is fully transparent and understandable to the user, as it returns a sequence of meaningful production rules that yield our model's result. Lastly, the generation process is invertible, so molecules can be translated from other popular representations such as SMILES. These superior properties make PolyGrammar more comprehensive and practical than existing representation schemes and generative models. Our full chemical design model (representation, generative model, and inverse model) is also efficient and straightforward to use in practice. For a polyurethane chain of length 20, the average generation time via PolyGrammar is 4 ms, and its translation from SMILES costs 11 ms on a PC with an Intel Core i7 CPU. The main contribution of the total generation time is from the context matching and rule selection at each production step, resulting in a linear escalation to the length of the chain. In order to generate a large number of different chains, one can easily use multi-processing techniques^[67] to generate numerous chains simultaneously. The overall generation time can be proportionally reduced by the number of parallels.

The current generative model of the PolyGrammar also only imitates the chain-growth polymerization. Although this polymerization mechanism has some benefits for the simulation of polyurethane chains,^[61] it would be ideal for our PolyGrammar to imitate step-growth polymerization as well. More advanced grammar such as universal grammar^[62] will be helpful to achieve this.

These aforementioned features are intriguing and will be implemented and demonstrated in future work. However, even without these augmentations, our proposed PolyGrammar takes an important step toward a more practical and comprehensive system for polymer discovery and exploration.

9. Conclusion

In summary, we propose a parametric context-sensitive grammar, called PolyGrammar, for the representation and generation of polymers. The recursive nature of grammar production enables the generation of any polymer chain using only a simple set of production rules. We also implement an algorithm that can transfer a SMILES string of a polymer chain to the sequence of production rules used to generate it. Capable of reproducing a large literature-collected dataset, this algorithm demonstrates the completeness and effectiveness of our PolyGrammar. Our PolyGrammar will benefit the polymer community in several ways. The most immediate contribution is our ability to efficiently generate an exhaustive collection of polymer samples. This corpus could be very powerful in conjunction with other methods (e.g., machine learning) to guide the synthesis of physical polymers and facilitate complex tasks like molecular discovery^[2–4] and property optimization.^[13,14,17] PolyGrammar is also helpful for the reverse engineering of polymer design and production. Our PolyGrammar serves as a blueprint to construct chemical design models for different classes of chemistries, including both organic and inorganic molecules. Eventually, PolyGrammar could improve chemical communication and exploration, by providing a more efficient and effective representation scheme that is widely suitable for complicated polymers.

Supporting Information

Supporting Information is available from the Wiley Online Library or from the author.

Acknowledgements

The authors would like to thank Dr. Xingcai Zhang and Dr. Ming Xiao from Harvard University, and Dr. Pengfei Zhang from Qingdao University for their helpful comments.

Conflict of Interest

The authors declare no conflict of interest.

Author Contributions

M.G. developed and implemented the algorithm and conducted the experiments. L.M. contributed to the organization and writing. W.S., T.E., and M.F. contributed to the development of production rules and polymer example collection. All authors edited and commented on the manuscript. W.M. initiated the original idea and supervised the research.

Data Availability Statement

All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supporting Information.

Keywords

context-sensitive grammar, generative model, polymer representation

Received: May 5, 2021

Revised: December 4, 2021

Published online:

- [1] E. J. Bjerrum, R. Threlfall, *Molecular generation with recurrent neural networks (RNNs)*, arXiv preprint arXiv:1705.04612 **2017**.
- [2] M. Olivecrona, T. Blaschke, O. Engkvist, H. Chen, *J. Cheminf.* **2017**, *9*, 48.
- [3] A. Gupta, A. T. Müller, B. J. Huisman, J. A. Fuchs, P. Schneider, G. Schneider, *Mol. Inf.* **2018**, *37*, 1700111.
- [4] M. Sumita, X. Yang, S. Ishihara, R. Tamura, K. Tsuda, *ACS Cent. Sci.* **2018**, *4*, 1126.
- [5] T. Blaschke, M. Olivecrona, O. Engkvist, J. Bajorath, H. Chen, *Mol. Inf.* **2018**, *37*, 1700123.
- [6] J. Lim, S. Ryu, J. W. Kim, W. Y. Kim, *J. Cheminf.* **2018**, *10*, 31.
- [7] W. Jin, R. Barzilay, T. Jaakkola, *Junction tree variational autoencoder for molecular graph generation*, arXiv preprint arXiv:1802.04364 **2018**.
- [8] P. B. Jørgensen, M. N. Schmidt, O. Winther, *Mol. Inf.* **2018**, *37*, 1700133.
- [9] E. Putin, A. Asadulaev, Y. Ivanenkov, V. Aladinskiy, B. Sanchez-Lengeling, A. Aspuru-Guzik, A. Zhavoronkov, *J. Chem. Inf. Model.* **2018**, *58*, 1194.
- [10] E. Putin, A. Asadulaev, Q. Vanhaelen, Y. Ivanenkov, A. V. Aladinskaya, A. Aliper, A. Zhavoronkov, *Mol. Pharmaceutics* **2018**, *15*, 4386.
- [11] O. Méndez-Lucio, B. Baillif, D.-A. Clevert, D. Rouquié, J. Wichard, *Nat. Commun.* **2020**, *11*, 10.

- [12] Ł. Maziarka, A. Pocha, J. Kaczmarczyk, K. Rataj, T. Danel, M. Warchoń, *J. Cheminf.* **2020**, *12*, 2.
- [13] F. Napolitano, Y. Zhao, V. M. Moreira, R. Tagliaferri, J. Kere, M. D'Amato, D. Greco, *J. Cheminf.* **2013**, *5*, 30.
- [14] G. Montavon, M. Rupp, V. Gobre, A. Vazquez-Mayagoitia, K. Hansen, A. Tkatchenko, K.-R. Müller, O. A. Von Lilienfeld, *New J. Phys.* **2013**, *15*, 095003.
- [15] Z. Wu, B. Ramsundar, E. N. Feinberg, J. Gomes, C. Geniesse, A. S. Pappu, K. Leswing, V. Pande, *Chem. Sci.* **2018**, *9*, 513.
- [16] C. W. Coley, W. Jin, L. Rogers, T. F. Jamison, T. S. Jaakkola, W. H. Green, R. Barzilay, K. F. Jensen, *Chem. Sci.* **2019**, *10*, 370.
- [17] H. Gao, T. J. Struble, C. W. Coley, Y. Wang, W. H. Green, K. F. Jensen, *ACS Cent. Sci.* **2018**, *4*, 1465.
- [18] J. S. Schreck, C. W. Coley, K. J. Bishop, *ACS Cent. Sci.* **2019**, *5*, 970.
- [19] B. Liu, B. Ramsundar, P. Kawthekar, J. Shi, J. Gomes, Q. Luu Nguyen, S. Ho, J. Sloane, P. Wender, V. Pande, *ACS Cent. Sci.* **2017**, *3*, 1103.
- [20] K. T. Butler, D. W. Davies, H. Cartwright, O. Isayev, A. Walsh, *Nature* **2018**, *559*, 547.
- [21] P. Raccuglia, K. C. Elbert, P. D. Adler, C. Falk, M. B. Wenny, A. Mollo, M. Zeller, S. A. Friedler, J. Schrier, A. J. Norquist, *Nature* **2016**, *533*, 73.
- [22] B. Barnes, B. Rice, A. Sifain, *Bull. Am. Phys. Soc.* **2020**, *65*, 1.
- [23] D. Fooshee, A. Mood, E. Gutman, M. Tavakoli, G. Urban, F. Liu, N. Huynh, D. Van Vranken, P. Baldi, *Mol. Syst. Des. Eng.* **2018**, *3*, 442.
- [24] P. Schwaller, T. Gaudin, D. Lanyi, C. Bekas, T. Laino, *Chem. Sci.* **2018**, *9*, 6091.
- [25] M. H. Segler, M. Preuss, M. P. Waller, *Nature* **2018**, *555*, 604.
- [26] A. B. Henson, P. S. Gromski, L. Cronin, *ACS Cent. Sci.* **2018**, *4*, 793.
- [27] L. M. Roch, F. Häse, C. Kreisbeck, T. Tamayo-Mendoza, L. P. Yunker, J. E. Hein, A. Aspuru-Guzik, *PLoS One* **2020**, *15*, 0229862.
- [28] D. Weininger, *J. Chem. Inf. Comput. Sci.* **1988**, *28*, 31.
- [29] S. Ash, M. A. Cline, R. W. Homer, T. Hurst, G. B. Smith, *J. Chem. Inf. Comput. Sci.* **1997**, *37*, 71.
- [30] J. J. Vollmer, *J. Chem. Educ.* **1983**, *60*, 192.
- [31] S. R. Heller, A. McNaught, I. Pletnev, S. Stein, D. Tchekhovskoi, *J. Cheminf.* **2015**, *7*, 23.
- [32] A. Drefahl, CurlySMILES: a chemical language to customize and annotate encodings of molecular and nanodevice structures. *J. Cheminf.* **2011**, *3*, 1.
- [33] T.-S. Lin, C. W. Coley, H. Mochigase, H. K. Beech, W. Wang, Z. Wang, E. Woods, S. L. Craig, J. A. Johnson, J. A. Kalow, K. F. Jensen, B. D. Olsen, *ACS Cent. Sci.* **2019**, *5*, 1523.
- [34] T. Zhang, H. Li, H. Xi, R. V. Stanton, S. H. Rotstein, *J. Chem. Inf. Modell.* **2012**, *10*, 2796.
- [35] H. Kajino, *Proc. Int. Conf. on Machine Learning*, IEEE, Piscataway, NJ **2019**, pp. 3183–3191.
- [36] W. Jin, R. Barzilay, T. Jaakkola, *Hierarchical Generation of Molecular Graphs using Structural Motifs*, arXiv preprint arXiv:2002.03230 **2020**.
- [37] L. H. Sperling, *Introduction to Physical Polymer Science*, John Wiley & Sons, New York **2005**.
- [38] Z. S. Petrovic, J. Ferguson, *Prog. Polym. Sci.* **1991**, *16*, 695.
- [39] J. M. G. Cowie, V. Arrighi, *Polymers: Chemistry and Physics of Modern Materials*, CRC Press, Boca Raton, FL **2007**.
- [40] D. C. Elton, Z. Boukouvalas, M. D. Fuge, P. W. Chung, *Mol. Syst. Des. Eng.* **2019**, *4*, 828.
- [41] D.-S. Cao, J.-C. Zhao, Y.-N. Yang, C.-X. Zhao, J. Yan, S. Liu, Q.-N. Hu, Q.-S. Xu, Y.-Z. Liang, *SAR QSAR Environ. Res.* **2012**, *23*, 141.
- [42] J. You, B. Liu, Z. Ying, V. Pande, J. Leskovec, *Proc. 32nd Int. Conf. on Neural Information Processing Systems*, **2018**.
- [43] Z. Zhou, S. Kearnes, L. Li, R. N. Zare, P. Riley, *Sci. Rep.* **2019**, *9*, 10752.
- [44] H. Dai, Y. Tian, B. Dai, S. Skiena, L. Song, *Syntax-directed variational autoencoder for structured data*, arXiv preprint arXiv:1802.08786 **2018**.
- [45] A. Nouira, N. Sokolovska, J.-C. Crivello, *Crystalgan: learning to discover crystallographic structures with generative adversarial networks*, arXiv preprint arXiv:1810.11203 **2018**.
- [46] Y. Tian, X. Zhang, H.-Z. Geng, H.-J. Yang, C. Li, S.-X. Da, X. Lu, J. Wang, S.-L. Jia, *RSC Adv.* **2017**, *7*, 53018.
- [47] G. Oertel, L. Abele, *Polyurethane Handbook: Chemistry, Raw Materials, Processing, Application, Properties*, Hanser Publishers, New York **1994**.
- [48] H.-W. Engels, H.-G. Pirkl, R. Albers, R. W. Albach, J. Krause, A. Hoffmann, H. Casselmann, J. Dormish, *Angew. Chem., Int. Ed.* **2013**, *52*, 9422.
- [49] Y. Niu, F. J. Stadler, T. He, X. Zhang, Y. Yu, S. Chen, *J. Mater. Chem. B* **2017**, *5*, 9477.
- [50] P. Czech, L. Okrasa, F. Méchin, G. Boiteux, J. Ulanski, *Polymer* **2006**, *47*, 7207.
- [51] Y. Li, O. Vinyals, C. Dyer, R. Pascanu, P. Battaglia, *Learning deep generative models of graphs*, arXiv preprint arXiv:1803.03324 **2018**.
- [52] R. Winter, F. Montanari, F. Noé, D.-A. Clevert, *Chem. Sci.* **2019**, *10*, 1692.
- [53] C. Berge, *Hypergraphs: Combinatorics of Finite Sets*, Vol. 45, Elsevier, New York **1984**.
- [54] J.-C. Bermond, M.-C. Heydemann, D. Sotteau, *Discrete Math.* **1977**, *18*, 235.
- [55] N. Chomsky, *IRE Trans. Info. Theory* **1956**, *2*, 113.
- [56] A. Lindenmayer, *J. Theor. Biol.* **1968**, *18*, 300.
- [57] P. Prusinkiewicz, A. Lindenmayer, *The Algorithmic Beauty of Plants*, Springer Science & Business Media, New York **2012**.
- [58] J. E. Hopcroft, R. Motwani, J. D. Ullman, *Acm Sigact News* **2001**, *32*, 60.
- [59] E. Gee, G. Liu, H. Hu, J. Wang, *Langmuir* **2018**, *34*, 10102.
- [60] S. S. Mahapatra, S. K. Yadav, H. J. Yoo, J. W. Cho, J.-S. Park, *Composites, Part B* **2013**, *45*, 165.
- [61] R. Ghoreishi, G. Suppes, *RSC Adv.* **2015**, *5*, 68361.
- [62] N. Chomsky, *Tool Module: Chomsky's Universal Grammar* **2018**.
- [63] D. E. Knuth, J. H. Morris, Jr, V. R. Pratt, *SIAM J. Comput.* **1977**, *6*, 323.
- [64] R. E. Korf, *Artificial Intelligence Search Algorithms*, ACM, New York **1999**.
- [65] J. Pearl, *Intelligent Search Strategies for Computer Problem Solving*, Addison Wesley, Boston **1984**.
- [66] M. Bhardwaj, S. Choudhury, S. Scherer, *Learning heuristic search via imitation*, arXiv preprint arXiv:1707.03034 **2017**.
- [67] J. Hunt, in *Multiprocessing, Advanced Guide to Python 3 Programming*, Springer, Cham **2019**, p. 363.
- [68] G. Rozenberg, *Handbook of Graph Grammars and Computing by Graph Transformation*, Vol. 1, World Scientific, Singapore **1997**.
- [69] H. D. Tran, C. Kim, L. Chen, A. Chandrasekaran, R. Batra, S. Venkatram, D. Kamal, J. P. Lightstone, R. Gurnani, P. Shetty, M. Ramprasad, J. Laws, M. Shelton, R. Ramprasad, *J. Appl. Phys.* **2020**, *128*, 171104.
- [70] C. Kim, A. Chandrasekaran, T. D. Huan, D. Das, R. Ramprasad, *J. Phys. Chem. C* **2018**, *122*, 31, 17575.
- [71] T. D. Huan, A. Mannodi-Kanakithodi, C. Kim, V. Sharma, G. Pilania, R. Ramprasad, *Sci. Data* **2016**, *3*, 160012.
- [72] O. Toshio, M. Yamazaki, "New stage of MatNavi, materials database at NIMS" **2012**.
- [73] W. C. Yount, D. M. Loveless, S. L. Craig, *J. Am. Chem. Soc.* **2005**, *127*, 14488.
- [74] M. Krenn, et al., *Machine Learning: Science and Technology 1.4* **2020**, 045024.
- [75] A. K. Nigam, et al., *Chemical Science* **2021**.
- [76] H. Ikebata, et al., *J. Computer-Aided Molecular Design* **31.4** **2017**, 379.