

Generative Foundation Model for On-demand Reverse Polymer Design

Haoke Qiu^{†,‡} and Zhao-Yan Sun^{*,†,‡}

[†]*State Key Laboratory of Polymer Physics and Chemistry, Changchun Institute of Applied Chemistry, Chinese Academy of Sciences, Changchun 130022, China*

[‡]*School of Applied Chemistry and Engineering, University of Science and Technology of China, Hefei 230026, China*

E-mail: zysun@ciac.ac.cn

Phone: +86 (0431) 85262896

Abstract

Forward screening and reverse design of drug molecules, inorganic molecules, and polymers with better properties are crucial engines for shortening the laboratory-to-market cycle. Particularly, due to the lack of large-scale datasets, polymer discovery based on materials informatics is more formidable. Despite this, polymer scientists have developed a series of machine learning models on polymer structure-property relationships using only small polymer datasets, thereby driving the forward screening process of polymers. However, the success of this paradigm ultimately hinges on the capacity of the candidate pool, while exhaustively enumerating all polymer structures through human imagination is challenging. Therefore, achieving on-demand reverse design of polymers is crucial. In this work, we curate a polymer dataset containing nearly one million polymeric structure-property pairs based on expert intuition. Using this dataset, we propose a generative pre-trained model for polymer on-demand generation using a large language model. The model produce polymers with 99.27% chemical

validity in top-1 generation mode (approximately 200k generated polymers), marking the highest reported success rate among polymer generative models. In addition, the average R^2 between the properties of the generated molecules and their expected values across 15 predefined properties is 0.96. To further assess the pre-trained model's performance in generating additional user-defined polymer properties for downstream tasks, we conduct fine-tuning experiments on three publicly available small polymer datasets using semi-template and template-free generation paradigm. Through these extensive experiments, we demonstrate that our pre-trained model and fine-tuned models are capable of achieving on-demand reverse design of polymers with specified properties, whether in (semi-)template generation or the more challenging template-free generation scenarios.

Introduction

On Earth, the candidate pool of potential materials is vast, with a reasonable estimate reaching 10^{60} .¹ However, this number could be even larger in reality, considering factors such as lattice defects in inorganic materials and the stochastic, multi-scale structures of polymers. In the monumental task of efficiently discovering candidate structures that can serve as 'materials', machine learning (ML) has demonstrated its impressive accuracy and efficiency on inorganic materials,^{2,3} metal materials,⁴ organic molecules,⁵ and polymer materials.⁶ Compared to other materials, the development of polymer materials is more challenging due to the limited data and complex cross-scale structure-property relationships.⁷⁻¹² In the implementation of ML-assisted polymeric materials discovery, two main types can be distinguished: **Forward Screening** and **Reverse Design**. The Forward Screening-based ML workflow relies on models to screen candidate structures from a predefined pool of candidate polymers. Various simple to complex ML models and deep learning models, such as convolutional neural networks,¹³ graph neural networks,^{14,15} recurrent neural networks,¹⁶ and the recent Transformer,^{12,17,18} have been used to establish surrogate models for polymer forward screen-

ing. Although this paradigm is highly effective, especially for small polymer datasets,^{10,19,20} there is a possibility that structures beyond human imagination may be overlooked.^{21,22} Fortunately, the latter paradigm achieves the direct on-demand reverse design of candidate structures that satisfy performance requirements, without the need for a predefined pool of candidates. This represents a more ideal strategy for on-demand polymer design and has the potential to generate candidate structures that experts may find difficult to imagine. At the technical level, this paradigm is implemented based on generative models and has seen groundbreaking and impressive cases in the design of small organic molecules,^{23–25} with the percentage of chemically valid small organic molecules generated exceeding 99%.²⁴

Inspired by these successes, polymer scientists are also attempting to develop generative models for polymers to meet evolving application demands. Batra et al.²¹ proposed a modified variational autoencoder (VAE) to generate polymer repeat units based on SMILES, but the chemical validity of the generated polymers was found to be less than 30%. Compared to generative models for organic small molecules, the significantly reduced chemical validity of generated polymers is mostly attributed to the presence of two special characters ('*') in polymer SMILES strings.^{11,21} These characters do not represent any chemical elements but denote unique polymerization points²⁶ and increase complexity of polymer generation,¹¹ leading to a sharp decrease in the performance of generative models trained on small polymer datasets. Indeed, training on larger datasets holds promise for enabling the model to learn this polymer knowledge. Meanwhile, polymer scientists are also attempting to represent polymers using another molecular representation method, namely molecular graphs, in order to enhance the chemical validity of generated polymers. Kim et al.,¹¹ Liu et al.,²⁷ and Gurnani et al.²² respectively trained polymer generative models using graph neural networks, resulting in an increase in the proportion of chemically valid molecules generated by the models to 16.07 to 89.40%, 44.03%, and 93%.

Although there has been improvement in the proportion of chemically valid polymers generated at present, current polymer generative models still face several challenges: 1)

First of all, there is a lack of pre-trained models for polymer generation. Due to the limitations of polymer datasets, polymer scientists usually train their polymer generative models from scratch using small datasets tailored to specific properties, such as dielectric performance.^{22,27} A unified pre-trained model holds the promise of leveraging small datasets for various polymer properties, enabling the accurate generation of polymers with multiple properties. 2) The current polymer generative models are trained based on SMILES-to-SMILES translation²¹ or graph-to-graph translation^{11,22,27} (or reconstruction), while the unsupervised nature of this strategy inherently requires more data to learn hidden chemical patterns. As a result, the ability of the current polymer generative models to generate chemically valid polymers is limited. 3) When generating new molecules, these unsupervised approaches typically rely on editing the numeric representation in the hidden layers of the generative models to generate new molecules via decoder layer. This implies that typically there needs to be a template polymer with the desired properties. However, finding such template polymers is actually a daunting task.²⁸ Besides, due to the large dimensionality of the editable numeric representation of the template polymer, the directionality of reverse design is partially out of control,^{23,27} introducing uncertainty to the screening task. 4) Additionally, as generation often occurs within the neighborhood of the template polymer,²⁷ it's difficult to efficiently explore the diverse polymer space as polymers with low structural similarity may exhibit similar properties. For instance, a polymer chain containing hydrogen bonding interactions may exhibit a similar glass transition temperature to another polymer chain containing multiple benzene rings.²⁸

To address these challenges, we refined a polymer structure-property dataset containing nearly 1,000,000 entries based on the largest unlabeled polymer dataset, PI1M.²⁹ Using this curated structure-property dataset, we propose a polymer generative pre-trained large language model (LLM) via supervised learning. The pre-trained model demonstrates impressive polymer generation capabilities, with chemical validity exceeding 99% when generating a total of approximately 200,000 polymers in top-1 mode. We evaluated the 15 input fun-

damental properties of the generated polymers, showing extremely high prediction accuracy with the expected values (the average R^2 is 0.96). We further tested its ability to generate polymers with other user-defined properties in three downstream tasks and also explored the feasibility of semi-template and completely template-free polymer generation. The results demonstrate its excellent performance in on-demand reverse polymer generation, and the generated polymers exhibits diverse structural features, which showcases the model's ability to thoroughly explore the polymer space.

Methods

Polymer structure-property dataset

The largest publicly available polymer structure-property dataset currently is PolyInfo, containing around 20,000 polymer structure-property pairs. However, this dataset is insufficient for training a LLM. Recently, a virtual polymer database, PI1M,²⁹ has been extended from PolyInfo, comprising nearly one million polymer structures but lacking corresponding property values. Researchers have utilized PI1M for unsupervised pretraining for polymer generation but this unsupervised pretraining paradigm results in the limited capacity of generating chemically valid polymers.²⁷ To achieve unprecedented large-scale supervised learning on this largest polymer dataset, we opted to compute foundational properties for each polymer structure in PI1M as descriptors. Due to the significant influence of molecular interactions and chain structure on the properties of polymers at the microscopic level, we carefully selected 15 descriptors related to the factors mentioned above. Specifically, we considered: 1) Molecular weight, providing an approximate constraint on the number and types of atoms in the repeat unit. 2) Hydrogen bonds, including the types and quantities of hydrogen bond donors and acceptors. 3) Atom types, including the number of heteroatoms apart from the common carbon and hydrogen atoms in polymers. 4) Chain structure, including the types and quantities of rings and the number of rotatable bonds, which account for the flexibility

of polymer molecules. For the specific list and the corresponding description, please refer to Section 1 of the Supporting Information (S1). The above foundational properties and their corresponding SMILES constitute our polymer structure-property dataset.

Prompt Engineering

Like any LLM, designing high-quality prompts is crucial for on-demand generation by the model. We computed 15 physicochemical properties for each polymer repeat unit's SMILES and concatenated them to form the input prompts for the pretrain model. Except for molecular weight, all other physicochemical properties are of integer type. To reduce the input token size and improve training efficiency, we also converted the molecular weight to an integer type. In fact, this approach, imprecisely specifying the molecular weight and instead slightly "fuzzifying" it, increases the model's freedom when generating new molecules, thus facilitating the generation of structurally diverse molecules (as shown in S2).

Model Settings

Currently, there are many open-source large language models (LLMs) available for pretraining in chemical tasks. However, our previous research demonstrates that LLM based on a deep understanding of chemical knowledge may perform better even on less data.¹⁸ Taking into account both model complexity and computational device requirements, we have chosen our previously developed PolyNC¹⁸ as the foundational model. PolyNC is a LLM based on polymer structures, capable of predicting various properties such as the glass transition temperature of polymers. This model possesses over 22 million parameters, thus providing ample capacity to accomplish this ambitious task.

Polymer generative pretraining via large-scale supervised learning

During the pretraining, the foundational properties will be concatenated as input, while the corresponding structures (SMILES) serve as the output (Figure 1a). Compared to generative models using SMILES-to-SMILES translation and graph-to-graph translation, our paradigm of property-to-SMILES aims to enable the model to capture more foundational properties of polymers and their corresponding structures (SMILES). We randomly partitioned the polymer structure-property dataset into a training set (80%, ~ 0.8 million) and a test set (20%, ~ 0.2 million) for pretraining. During fine-tuning for "semi-template" generation, we added other polymer's properties as additional vectors to the input prompt (Figure 1(b)-(d)).

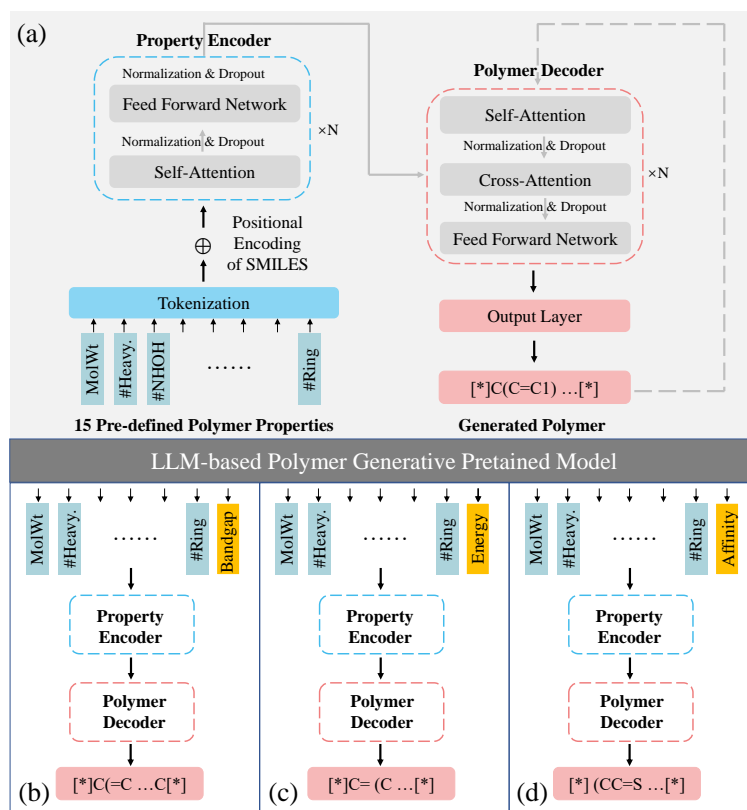


Figure 1: Generation workflow. (a) Model Pretraining. Using a set of 15 predefined structural features related to polymers as input, supervised learning for conditional generation has demonstrated impressive accuracy in generating polymer repeat units that satisfy these features, expressed as SMILES strings. Subsequently, we validated the model's ability to generate polymers with specific properties by assessing its capability to generate polymers with specified band gaps (b), atomization energy(c), and electron affinity(d).

Model Metrics

For molecular generative models, the primary metric of interest is the chemical **Validity** (the percentage of chemically valid molecules), which is our foremost consideration. Additionally, LLM can produce multiple outputs for the same input (i.e., top-k generation), which is beneficial for generating structurally diverse candidate molecules and helps assess the stability of the model's generation capability. Therefore, for top-k generation scenarios, we also evaluate the **Uniqueness** (the percentage of chemically valid molecules generated that are mutually unique in each generation of the k times generation) and **Novelty** (the percentage of generated valid molecules not in the training set and the test set in each generation of the k times generation) of the generated polymers.

It is worth mentioning that previous polymer generative models rarely discussed the similarity of generated polymers to existing polymers and, more importantly, the synthetic feasibility of the generated polymers. Thus, in top-k generation scenarios, we additionally assess the **Similarity** (i.e., Tanimoto similarity³⁰) and synthetic feasibility³¹ (**SA**score) of the generated polymers.

Results and discussion

Performance of the pre-trained model on polymer generation

Top-1 generation

We first ran top-1 generation testing (i.e., generating one polymer for each input) using the pre-trained model on the test set, resulting in a total of 199,159 samples generated. Among the generated polymers, 99.27% were chemically valid molecules, which represents the highest value among existing polymer generative models to date and was achieved on the largest test sample (**Table 1**). This demonstrates that the model has deeply learned the mapping between polymeric fundamental properties and SMILES after pretraining via

large-scale supervised learning.

Table 1: Performance of the pre-trained model on the test test via top-1 generation. # Data: number of training data; # Gen.: number of generated polymers; Val.: validity; UNC: unconditional; COND: conditional. ^a Results for two polymer properties (glass transition temperature and band gap). ^b Results for one polymer properties (logP). ^c Results for 15 polymer properties (as illustrated in Methods).

Model	Architecture	Mode	# Data	# Gen.	Val./% \uparrow	Average R^2 \uparrow
SD-VAE ²¹	CNN	UNC	250k	1k	13-27	0.65 ^a
polyG2G ²²	GNN	UNC	13k	58k	93	
IGGM ²⁷	GNN	UNC	250k	10k	44.03	
Mole. Chef ¹¹	GNN	UNC	120k		16.07-89.40	0.96 ^b
Ours	Transformer	CND	800k	199k	99.27	0.96^c

We conducted a statistical analysis of the types of chemical elements present in our polymer structure-property dataset and the generated polymers (Figure 2). The most abundant elements are C, N, and O, followed by other inorganic elements such as S and F, while metal elements constitute a smaller proportion, which aligns with the empirical knowledge in polymer science. The distribution of element proportions in the training set (Figure 2a) is similar to that in the test set (Figure 2b), indicating a relatively uniform dataset partition. Interestingly, in the polymers generated using the top-1 mode, some metal elements are not generated (Figure 2c). This is because that each LLM generates tokens (i.e., chemical elements) based on the probability of each token’s occurrence, so tokens with very low probabilities may not be generated. If necessary, this can be optimized by increasing the proportion of metal polymers in downstream tasks.

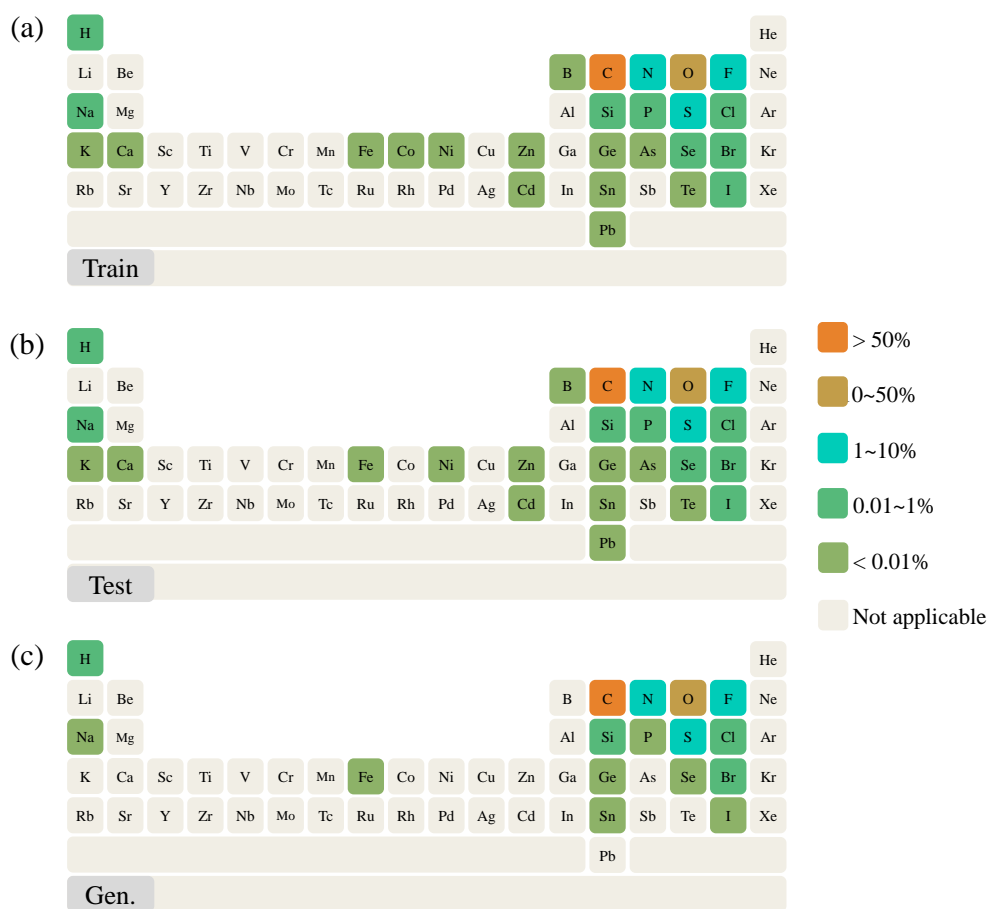


Figure 2: Statistical analysis of the types of chemical elements of the: training set (a), test set (b) and generated polymers (c). Gen.: Generated.

In order to demonstrate whether the generated polymers possess the expected fundamental properties specified in the input parameters, we examined the aforementioned properties of the generated polymers, where we found a high degree of agreement between them (Figure 3). For the chemically valid and unique polymers generated by the pre-trained model, the average R^2 value across the 15 polymer properties is 0.96. This indicates that the model can preliminarily achieve the on-demand generation of polymers with specified properties. This result instills confidence in our model's ability to generate polymers with other properties, as discussed in the *Applications of the pre-trained model in generating polymers with specific properties* section. Interestingly, for certain properties, such as the number of aliphatic carboncycles (NumAliphaticCarboncycles) and the number of aliphatic heterocycles (Nu-

mAliphaticHeterocycles), the generated polymers exhibit relatively poor consistency. However, these inconsistencies are advantageous for the model to produce structurally diverse polymers, facilitating the the generation of various ring structures (as illustrated in **S2**).

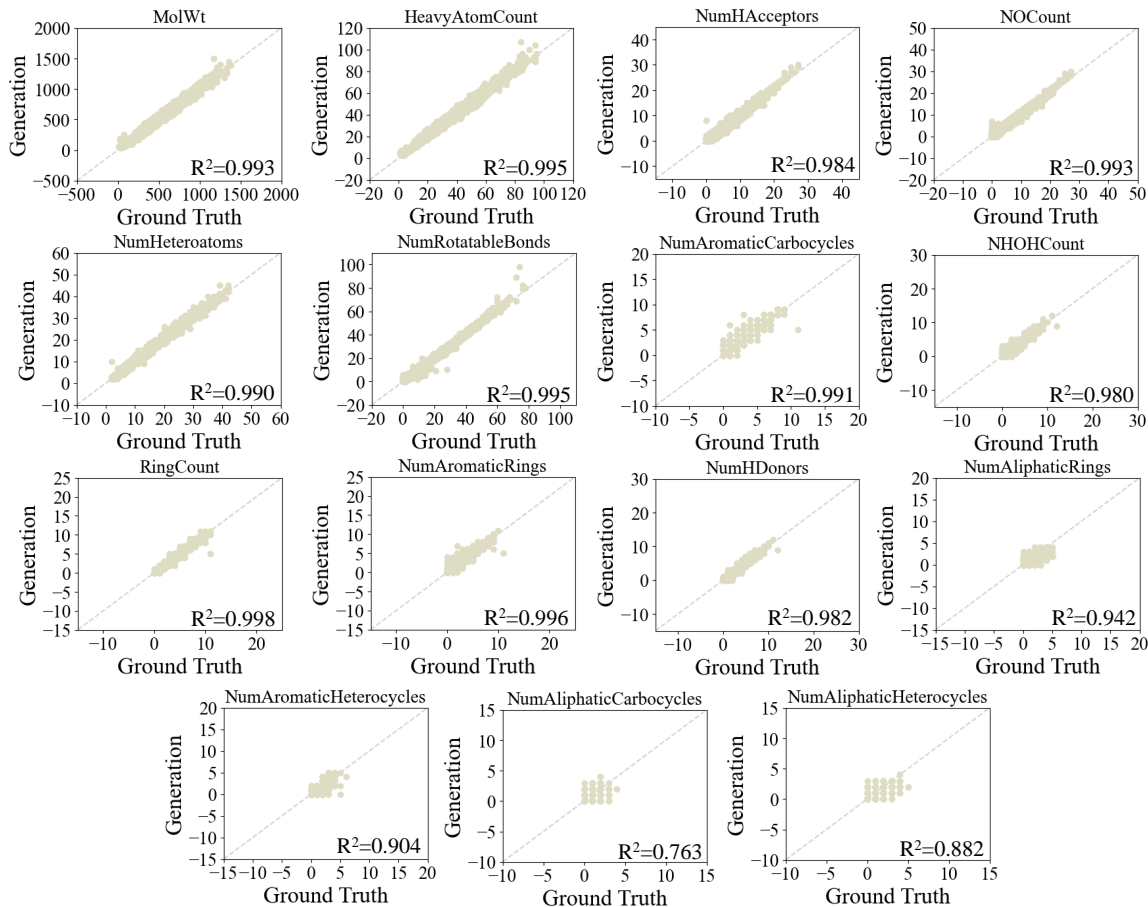


Figure 3: The fitting plots of the properties generated by the model (Generation) against the expected properties (Ground Truth).

Top-k ($k > 1$) generation

Due to the stochastic and probabilistic nature of generation of LLMs, to validate the stability of the pre-trained model's generation capability, we examined the model's performance in top-k generation. Specifically, we generated three (top-3), five (top-5), and ten (top-10) samples for the same input, evaluating the Validity, Uniqueness, and Novelty of the generated polymers. Additionally, we assessed the Tanimoto similarity coefficient between the generated polymers and input samples, as well as the synthesizability (SAscore) of the

generated polymers. The results are reflected as **Table 2**. In the top-k generation mode, the model tends to be "adventurous" in its generation, resulting in polymers that maintain a high level of uniqueness and novelty. Compared to top-1 generation (Figure 4a), top-k generation expands the chemical space of generated polymers, even extending beyond the chemical space corresponding to the training set (Figure 4b). Regardless of the value of k, the chemical similarity between the generated polymer and the polymer corresponding to the input prompt is consistently low. In contrast, generating new molecules based on artificially modified latent representations of molecules tend to produce molecules with very high similarity.²⁷ However, this "adventurous" generation can also lead to the generation of chemically invalid polymers, resulting in a slight decrease in the chemical validity of the generated polymers (above 97%), but still higher than previous polymer generative models. Impressively, for all valid, unique and novel molecules generated under the top-10 mode (totaling 1,828,027), their synthesizability did not become more challenging, demonstrating a synthesizability similar to that of PI1M as illustrated in Figure 4c.

Table 2: Performance of the pre-trained model on the test test via top-k (k=3,5,10) generation.

Metric	Top-3	Top-5	Top-10
Validity \uparrow	97.75 \pm 0.0001	97.76 \pm 0.0002	97.75 \pm 0.0004
Uniqueness \uparrow	99.07 \pm 0.0001	99.06 \pm 0.0002	99.08 \pm 0.0001
Novelty \uparrow	93.56 \pm 0.0009	93.73 \pm 0.0005	94.01 \pm 0.0006
Similarity \downarrow	0.302 \pm 0.0002	0.303 \pm 0.0002	0.306 \pm 0.0002
SAscore \downarrow	3.84 \pm 0.77	3.83 \pm 0.77	3.85 \pm 0.77

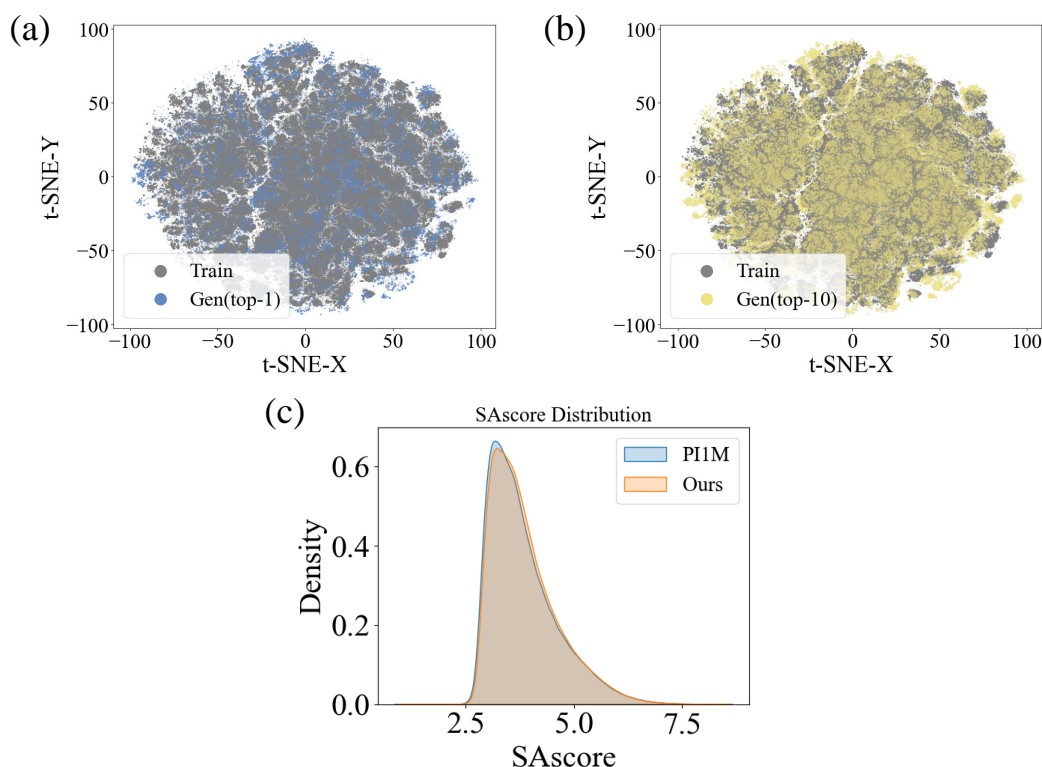


Figure 4: The chemical space representation of polymers after t-SNE dimensionality reduction, with the background color indicating the randomly selected training set from PI1M and (a) represents top-1 generation, with (b) representing top-10 generation. (c) illustrates the SAScore of PI1M and the valid, unique, and novel molecules generated under the top-10 mode.

Applications of the pre-trained model in generating polymers with specific properties

The above results demonstrates the impressive capability of the pre-trained model in generating polymers with specified foundational properties. Then we further assessed the model's ability to generate polymers with other specified properties.

In principle, polymer generation based on SMILES-to-SMILES translation or graph-to-graph translation require a template polymer that meets the desired performance.^{11,21,22,27} By editing the latent representation of the template polymer and decoding this representation, new polymers can be generated. We refer to this paradigm as **template-based** polymer generation. However, due to the randomness in editing the latent representation, achieving

on-demand design through this method is limited. Additionally, finding template polymers that meet the requirements is also a challenging task. Instinctively, By incorporating the target properties directly into the input prompts in a similar manner to the pre-training phase, there is potential to achieve on-demand generation of polymers without providing template polymers. Though this approach represents an advancement compared to template-based polymer generation, the input at this stage includes not only the target properties but also the 15 fundamental polymer properties we defined. Therefore we define this paradigm as **semi-template** generation.

Then, we tested the performance of the pre-trained model in the semi-template generation scenario. We finetuned this pre-trained model on ten public polymers datasets of different properties (**S3**) to obtain expert LLMs for each property. However, since experimentally validating the properties of generated polymers on a large scale is resource-intensive, we attempted to train proxy models for each dataset to efficiently validate the properties of the generated polymers on large-scale. To be specific, we utilized graph neural networks, known for their excellent performance in molecular property prediction, to train proxy models for each property. Since the accuracy, i.e., coefficient of determination (R^2), of the proxy models is crucial in assessing the properties of the generated polymers and the performance of expert LLMs, we selected the top three proxy models with the highest R^2 (exceeding 0.9, detailed in S3) and the corresponding dataset as subsequent case studies. These proxy models are tailored for the following polymer properties: band gap, atomization energy, and electron affinity. We then finetuned the pre-trained model using these datasets individually, with the same data partitioning as when training the respective proxy models. Each polymer property served as an additional vector added to the input prompt (see Figure 1 (b)-(d)). After fine-tuning, the loss of the expert LLMs exhibited convergence (Figure 5a, 6a and 7a).

In the generation of expert LLMs, we conducted top-5 generation (i.e., generating five samples for each input) and repeated the process for three rounds, to mitigate the randomness of LLM and assess the model's ability to generate multiple polymers satisfying the target

properties. Here, our main focus is to investigate the feasibility and reliability of LLM-based polymer generative model in practical usage. Therefore, we do not delve into detailed discussions regarding the uniqueness of generated polymers in the following sections.

We selected samples from the test sets for each property, following the criterion that the proxy model's prediction for the sample closely matches the ground truth, aiming to enhance the rationality of utilizing the proxy model to assess the properties of generated polymers (i.e., at least, the proxy model should be sufficiently accurate in predicting the properties of input samples). Moreover, we also aimed to ensure that the selected samples exhibit outstanding properties whenever possible.

Band gap The band gap of polymer holds significant importance in the advancement of polymer-based electronic and photonic devices, as it profoundly impacts their functionality across domains such as organic photovoltaics and light-emitting diodes. We take the example of generating polymers with wide band gaps (greater than 6 eV).²² Following the aforementioned selection criteria, we opted for the structure depicted in Figure 5a as the input sample, with its 15 fundamental properties plus its band gap serves as the input prompt. Across 3 rounds of top-5 generation, the expert LLM on band gap yielded 14 out of 15 chemically valid polymers, with 13 out of 14 showcasing novel structures from the input sample. Impressively, the predictions of the proxy model showcased that these 13 novel samples demonstrate properties that align with the target band gap (with a margin of error of 5%).

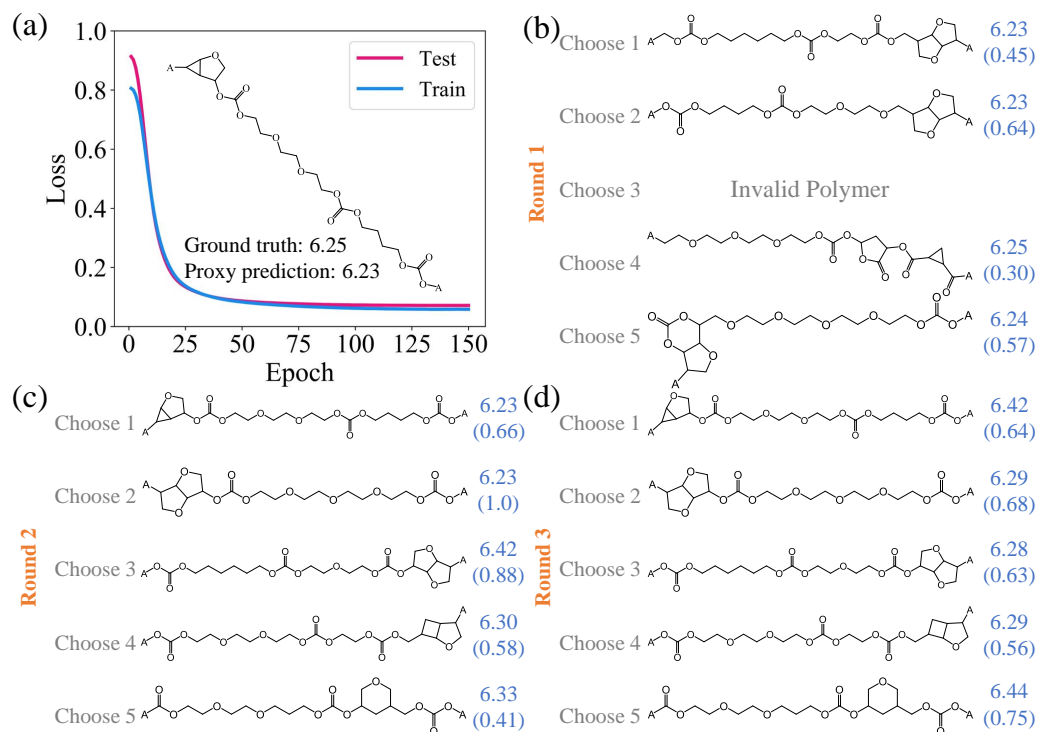


Figure 5: Loss of the expert LLM on band gap and the input sample (a) and on-demand inverse generation on band gap task (b)-(d). The blue numbers in subfigures (b)-(d) represent the band gaps predicted by the proxy model (relative to the similarity with the input sample), i.e., *predicted band gap(similarity)*. During the fine-tuning of this task, the training set consisted of 3042 samples, while the test set comprised 338 samples.

Atomization Energy The atomization energy of polymers reflects the strength and stability of the bonds within polymer molecules. Similarly, from the test set of the atomization energy database, we selected a polymer with high atomization energy as the input sample (Figure 6a). The 15 fundamental properties plus the atomization energy of this polymer is used as the input prompt. After three rounds of top-5 generation, the expert LLM on atomization energy generated 100% chemically valid and novel polymers. It is noteworthy that the data size of this dataset is too small for a LLM, resulting in a slight decrease in the accuracy of generated polymers. According to the predictions of the proxy model, 11 out of 15 polymers exhibit properties that align with the target atomization energy (with a margin of error of 5%).




Figure 6: Loss of the expert LLM on atomization energy and the input sample (a) and on-demand inverse generation on atomization energy task (b)-(d). The blue numbers in subfigures (b)-(d) represent the band gaps predicted by the proxy model (relative to the similarity with the input sample), i.e., *predicted atomization energy(similarity)*. During the fine-tuning of this task, the training set consisted of 351 samples, while the test set comprised 39 samples.

Electron Affinity The electron affinity of polymers reflects the polymer molecule's ability to accept electrons, a property crucial in photovoltaic applications and other electronic applications of polymers. Unfortunately, datasets for this property are also relatively small. We chose the 15 fundamental properties plus the electron affinity of the structure depicted in Figure 7a as the input prompt. From the results of three rounds of top-5 generation, 100% of the generated polymers were chemically valid, with 14 out of 15 being novel. According to predictions from the proxy model, 10 out of 14 polymers exhibit properties that align with the target electron affinity (with a margin of error of 5%).

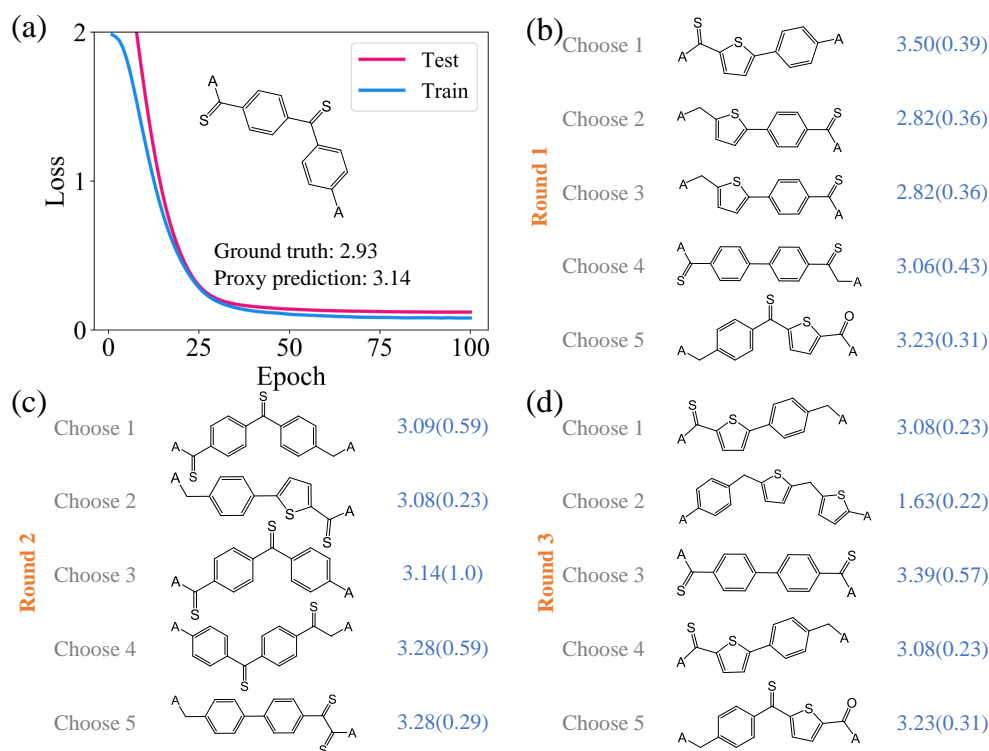


Figure 7: Loss of the expert LLM on electron affinity and the input sample (a) and on-demand inverse generation on electron affinity task (b)-(d). The blue numbers in subfigures (b)-(d) represent the band gaps predicted by the proxy model (relative to the similarity with the input sample), i.e., *predicted electron affinity(similarity)*. During the fine-tuning of this task, the training set consisted of 331 samples, while the test set comprised 37 samples.

Template-free generation: an ambitious task towards polymer on-demand design

Compared to previous template-based approaches,^{11,21,22,27} the semi-template method introduced above takes target properties as part of input, enabling the generation of polymers with specified properties. This represents an advancement over entirely template-based polymer generation. One more challenging goal is to achieve template-free polymer generation. The potential scenario for this paradigm is to provide the generative model with only a desired property value, allowing the model to freely generate structures that meet the requirements. Clearly, this design paradigm is ambitious yet more challenging due to the contradiction between the infinite chemical space of polymers and the limited training data available. To

assess the feasibility of our pre-trained model in this challenging task, we conducted fine-tuning tests using the band gap dataset due to its larger number of data entries. During the implement, we utilized only the value of band gap as input, with the corresponding polymer repeat unit SMILES as output (Figure 8a).

Meanwhile, in order to simultaneously achieve a meaningful objective, our aim is to have the fine-tuned expert LLM generate polymers with higher band gaps. We used 6.5 eV (higher than previous semi-template generation task) as target, then the model was expected to generate polymer structures with band gaps around 6.5 eV. Throughout the generation process, we continued to utilize the top-5 generation mode and repeated the process for three rounds to assess the stability of the LLM-based polymer generative model. The generation results in Figure 8b indicate that the proportion of chemically valid molecules generated is 100%, which is a significant prerequisite for the success of this task. Furthermore, in this novel task previously unexplored by polymer scientists, as verified by the proxy model, the expert LLM can produce no fewer than 2 samples with target band gap (with a margin of error of 5%) in each round of generation (marked in red font in Figure 8b). In total, 9 polymers exhibit properties that align with the target band gap (with a margin of error of 5%). Interestingly, compared to template-based and semi-template-based polymer generation methods mentioned earlier, the template-free approach generates a more diverse range of polymer structures, showcasing the model's freedom to explore the polymer space. These results demonstrate the feasibility of on-demand reverse polymer generation without template.




Figure 8: (a) Finetuning the pre-trained model for the temple-free generation. (b) The generation results via temple-free generation. The model will target the given value and generate polymer structures with values around the specified threshold.

Discussion

In this work, we proposed a generative pre-trained model based on LLMs for on-demand reverse polymer design. The pre-trained model was trained on a meticulously curated dataset containing nearly one million polymer structures and fundamental properties, crafted based on expert knowledge. Evaluation on a test set of nearly 200,000 samples revealed that the model generated chemically valid molecules with a proportion of 99.27%. Through further top-10 generation, the pre-trained model designed over 1.8 million valid and novel polymer structures, effectively doubling the entries of the off-the-shelf polymer datasets. These data, along with widely known datasets like PI1M, can offer a richer candidate pool for paradigms based on forward screening. To achieve the generation of polymers with other specific properties, we fine-tuned the pre-trained model on three publicly available polymer

property datasets, resulting in expert LLMs tailored to each property. The generation results of these expert LLMs demonstrate the powerful capability of the model in on-demand reverse generation. However, for more precise on-demand design, we advocate for greater efforts from the polymer community to expand polymer property datasets. Additionally, we attempted an ambitious task using the dataset with a relatively large amount of data on band gap, aiming for completely template-free polymer generation. The results indicate that the fine-tuned expert LLM can achieve on-demand reverse polymer generation based solely on the provided values of the desired polymer properties.

In summary, we have demonstrated a pre-trained model for on-demand reverse generation of polymers, and its performance on multiple downstream datasets indicates its broad applicability and transferability. Meanwhile, by employing more advanced polymer representations, such as BigSMILES,³² coupled with a larger amount of polymer data, there is potential to further enhance the model's performance in on-demand polymer generation. However, this requires ongoing efforts, including but not limited to advancing the acquisition of large-scale BigSMILES strings and collecting polymer multimodal, multiscale data.

Acknowledgement

We thank the support from the National Key R&D Program of China (No. **2022YFB3707303**), and the National Natural Science Foundation of China (No. **21833008** and **52293471**). We sincerely thank Prof. Rampi Ramprasad, Prof. Ryo Yoshida, Prof. Ying Li, Prof. Zhi-Min Dang, Prof. Nicholas E. Jackson, Prof. Arash Khajeh and their colleagues for their pioneering works in the field of polymer generation models. We are also grateful to the Network and Computing Center in Changchun Institute of Applied Chemistry for the hardware support.

Supporting Information Available

The following files are available free of charge.

- Polymer physicochemical properties selected
- Structurally diverse polymers generation
- Proxy model for evaluation the properties of the generated polymers

References

- (1) Sanchez-Lengeling, B.; Aspuru-Guzik, A. Inverse molecular design using machine learning: Generative models for matter engineering. *Science* **2018**, *361*, 360–365.
- (2) Merchant, A.; Batzner, S.; Schoenholz, S. S.; Aykol, M.; Cheon, G.; Cubuk, E. D. Scaling Deep Learning for Materials Discovery. *Nature* **2023**, *624*, 80–85.
- (3) Szymanski, N. J. et al. An Autonomous Laboratory for the Accelerated Synthesis of Novel Materials. *Nature* **2023**, *624*, 86–91.
- (4) Nandy, A.; Duan, C.; Taylor, M. G.; Liu, F.; Steeves, A. H.; Kulik, H. J. Computational Discovery of Transition-metal Complexes: From High-throughput Screening to Machine Learning. *Chemical Reviews* **2021**, *121*, 9927–10000.
- (5) Li, X.; Maffettone, P. M.; Che, Y.; Liu, T.; Chen, L.; Cooper, A. I. Combining machine learning and high-throughput experimentation to discover photocatalytically active organic molecules. *Chem. Sci.* **2021**, *12*, 10742–10754.
- (6) Lu, H.; Diaz, D. J.; Czarnecki, N. J.; Zhu, C.; Kim, W.; Shroff, R.; Acosta, D. J.; Alexander, B. R.; Cole, H. O.; Zhang, Y.; Lynd, N. A.; Ellington, A. D.; Alper, H. S. Machine Learning-Aided Engineering of Hydrolases for PET Depolymerization. *Nature* **2022**, *604*, 662–667.

- (7) Wu, S.; Kondo, Y.; Kakimoto, M.-a.; Yang, B.; Yamada, H.; Kuwajima, I.; Lambert, G.; Hongo, K.; Xu, Y.; Shiomi, J.; Schick, C.; Morikawa, J.; Yoshida, R. Machine-Learning-Assisted Discovery of Polymers with High Thermal Conductivity Using a Molecular Design Algorithm. *npj Comput Mater* **2019**, *5*, 66.
- (8) Kuenneth, C.; Rajan, A. C.; Tran, H.; Chen, L.; Kim, C.; Ramprasad, R. Polymer Informatics with Multi-Task Learning. *Patterns* **2021**, *2*, 100238.
- (9) Yang, J.; Tao, L.; He, J.; McCutcheon, J. R.; Li, Y. Machine Learning Enables Interpretable Discovery of Innovative Polymers for Gas Separation Membranes. *Sci. Adv.* **2022**, *8*, eabn9545.
- (10) Barnett, J. W.; Bilchak, C. R.; Wang, Y.; Benicewicz, B. C.; Murdock, L. A.; Bereau, T.; Sanat K. Kumar, Designing Exceptional Gas-Separation Polymer Membranes Using Machine Learning. *Science Advances* **2020**, *6*, eaaz4301.
- (11) Kim, S.; Schroeder, C. M.; Jackson, N. E. Open Macromolecular Genome: Generative Design of Synthetically Accessible Polymers. *ACS Polymers Au* **2023**, *3*, 318–330.
- (12) Xu, C.; Wang, Y.; Barati Farimani, A. TransPolymer: A Transformer-based Language Model for Polymer Property Predictions. *npj Comput Mater* **2023**, *9*, 64.
- (13) Yan, C.; Feng, X.; Li, G. From Drug Molecules to Thermoset Shape Memory Polymers: A Machine Learning Approach. *ACS Applied Materials & Interfaces* **2021**, *13*, 60508–60521, PMID: 34878247.
- (14) Queen, O.; McCarver, G. A.; Thatigotla, S.; Abolins, B. P.; Brown, C. L.; Maroulas, V.; Vogiatzis, K. D. Polymer Graph Neural Networks for Multitask Property Learning. *npj Comput Mater* **2023**, *9*, 90.
- (15) Aldeghi, M.; Coley, C. W. A Graph Representation of Molecular Ensembles for Polymer Property Prediction. *Chem. Sci.* **2022**, *13*, 10486–10498.

- (16) Simine, L.; Allen, T. C.; Rossky, P. J. Predicting optical spectra for optoelectronic polymers using coarse-grained models and recurrent neural networks. *Proceedings of the National Academy of Sciences* **2020**, *117*, 13945–13948.
- (17) Kuenneth, C.; Ramprasad, R. polyBERT: A Chemical Language Model to Enable Fully Machine-Driven Ultrafast Polymer Informatics. *Nat Commun* **2023**, *14*, 4099.
- (18) Qiu, H.; Liu, L.; Qiu, X.; Dai, X.; Ji, X.; Sun, Z.-Y. PolyNC: A Natural and Chemical Language Model for the Prediction of Unified Polymer Properties. *Chem. Sci.* **2024**, *15*, 534–544.
- (19) Bradford, G.; Lopez, J.; Ruza, J.; Stolberg, M. A.; Osterude, R.; Johnson, J. A.; Gomez-Bombarelli, R.; Shao-Horn, Y. Chemistry-Informed Machine Learning for Polymer Electrolyte Discovery. *ACS Cent. Sci.* **2023**, *9*, 206–216.
- (20) Qiu, H.; Qiu, X.; Dai, X.; Sun, Z.-Y. Design of Polyimides with Targeted Glass Transition Temperature Using a Graph Neural Network. *J. Mater. Chem. C* **2023**, *11*, 2930–2940.
- (21) Batra, R.; Dai, H.; Huan, T. D.; Chen, L.; Kim, C.; Gutekunst, W. R.; Song, L.; Ramprasad, R. Polymers for Extreme Conditions Designed Using Syntax-Directed Variational Autoencoders. *Chem. Mater.* **2020**, *32*, 10489–10500.
- (22) Gurnani, R.; Kamal, D.; Tran, H.; Sahu, H.; Scharm, K.; Ashraf, U.; Ramprasad, R. polyG2G: A Novel Machine Learning Algorithm Applied to the Generative Design of Polymer Dielectrics. *Chem. Mater.* **2021**, *33*, 7008–7016.
- (23) Gómez-Bombarelli, R.; Wei, J. N.; Duvenaud, D.; Hernández-Lobato, J. M.; Sánchez-Lengeling, B.; Sheberla, D.; Aguilera-Iparraguirre, J.; Hirzel, T. D.; Adams, R. P.; Aspuru-Guzik, A. Automatic Chemical Design Using a Data-Driven Continuous Representation of Molecules. *ACS Central Science* **2018**, *4*, 268–276, PMID: 29532027.

- (24) Weiss, T.; Mayo Yanes, E.; Chakraborty, S.; Cosmo, L.; Bronstein, A. M.; Gershoni-Poranne, R. Guided Diffusion for Inverse Molecular Design. *Nature Computational Science* **2023**, *3*, 873–882.
- (25) Swanson, K.; Liu, G.; Catacutan, D. B.; Arnold, A.; Zou, J.; Stokes, J. M. Generative AI for Designing and Validating Easily Synthesizable and Structurally Novel Antibiotics. *Nature Machine Intelligence* **2024**, *6*, 338–353.
- (26) Ma, R.; Zhang, H.; Luo, T. Exploring High Thermal Conductivity Amorphous Polymers Using Reinforcement Learning. *ACS Appl. Mater. Interfaces* **2022**, *14*, 15587–15598.
- (27) Liu, D.-F.; Zhang, Y.-X.; Dong, W.-Z.; Feng, Q.-K.; Zhong, S.-L.; Dang, Z.-M. High-Temperature Polymer Dielectrics Designed Using an Invertible Molecular Graph Generative Model. *Journal of Chemical Information and Modeling* **2023**, *63*, 7669–7675, PMID: 38061777.
- (28) Qiu, H.; Wang, J.; Qiu, X.; Dai, X.; Sun, Z.-Y. Heat-Resistant Polymer Discovery by Utilizing Interpretable Graph Neural Network with Small Data. *Macromolecules* **2024**, *57*, 3515–3528.
- (29) Ma, R.; Luo, T. PI1M: A Benchmark Database for Polymer Informatics. *J. Chem. Inf. Model.* **2020**, *60*, 4684–4690.
- (30) Bajusz, D.; Rácz, A.; Héberger, K. Why is Tanimoto index an appropriate choice for fingerprint-based similarity calculations? *J. Cheminform.* **2015**, *7*, 20.
- (31) Ertl, P.; Schuffenhauer, A. Estimation of Synthetic Accessibility Score of Drug-like Molecules Based on Molecular Complexity and Fragment Contributions. *J. Cheminform.* **2009**, *1*, 8.
- (32) Lin, T.-S.; Coley, C. W.; Mochigase, H.; Beech, H. K.; Wang, W.; Wang, Z.; Woods, E.; Craig, S. L.; Johnson, J. A.; Kalow, J. A.; Jensen, K. F.; Olsen, B. D. BigSMILES: A

Structurally-Based Line Notation for Describing Macromolecules. *ACS Central Science*
2019, 5, 1523–1531.