

Benchmarking Study of Deep Generative Models for Inverse Polymer Design

Tianle Yue,¹ Lei Tao,² Vikas Varshney,³ Ying Li^{1*}

¹Department of Mechanical Engineering, University of Wisconsin-Madison, Madison, WI 53706, United States

²Department of Mechanical Engineering, University of Connecticut, Storrs, CT 06269, United States

³Materials and Manufacturing Directorate, Air Force Research Laboratory, Wright-Patterson Air Force Base, OH 45433, United States

*Corresponding author: yli2562@wisc.edu

SUMMARY

Molecular generative models based on deep learning have increasingly gained attention for their ability in *de novo* polymer design. However, there remains a knowledge gap in the thorough evaluation of these models. This benchmark study explores *de novo* polymer design using six popular deep generative models: Variational Autoencoder (VAE), Adversarial Autoencoder (AAE), Objective-Reinforced Generative Adversarial Networks (ORGAN), Character-level Recurrent Neural Network (CharRNN), REINVENT, and GraphINVENT. Various metrics highlighted the excellent performance of CharRNN, REINVENT, and GraphINVENT, particularly when applied to the real polymer dataset, while VAE and AAE show more advantages in generating hypothetical polymers. The CharRNN, REINVENT, and GraphINVENT models were further trained on real polymers utilizing reinforcement learning methods, targeting the generation of hypothetical polymers with high glass transition temperatures. The findings of this study provide critical insights into the capabilities and limitations of each generative model, offering valuable guidance for future endeavors in polymer design and discovery.

Keywords: polymer design, generative model, reinforcement learning, glass transition temperature

1. INTRODUCTION

Polymers represent an important class of materials, known for their exceptional versatility in numerous properties, including thermal, mechanical, optical, and dielectric characteristics.¹⁻⁵ Plentiful studies are recently dedicated to the molecular design of new polymers endowed with exceptional properties.⁶⁻¹¹ With the recent advancements in deep learning and its application in polymer science and engineering, *de novo* polymer design has been recognized as a promising method to expedite the design and discovery of new high-performance polymer materials.^{7, 12-14}

Extensive research has been conducted on the *de novo* design of polymers, with researchers adopting various approaches, especially for proposing new hypothetical polymer structures, as summarized in **Figure 1**. For example, Sharma et al. employed a polymer building block approach and high-throughput density functional theory (DFT) to design organic polymers with high energy storage capabilities.¹⁵ Initially, repeat units were created using four building blocks within each unit, with each block selected from a pool comprising -CH₂-, -C₆H₄-, -C₄H₂S-, -NH-, -CO-, -O-, and -CS-. These blocks were chosen due to their prevalence in polymer backbones. This was followed by a multi-stage screening process involving quantum mechanics-based searches and molecular dynamics techniques. The final phase included synthesizing and testing the most promising polymers, validating this approach for material selection. Similarly, Li et al. devised novel polysulfates by leveraging their knowledge of known polymer structures and the characteristics of functional groups.¹⁶ They then confirmed these structures' high T_g and band gap (E_g) values through experimental synthesis and characterizations. The advantages of these two studies are that they allow for control over the structural complexity of the hypothetical polymers and enable the prediction of their overall properties based on the characteristics of functional groups or substructures. However, such a combination method of polymer building blocks becomes quite challenging when there is a desire to obtain a large number of candidates, in particular, on the order of millions.

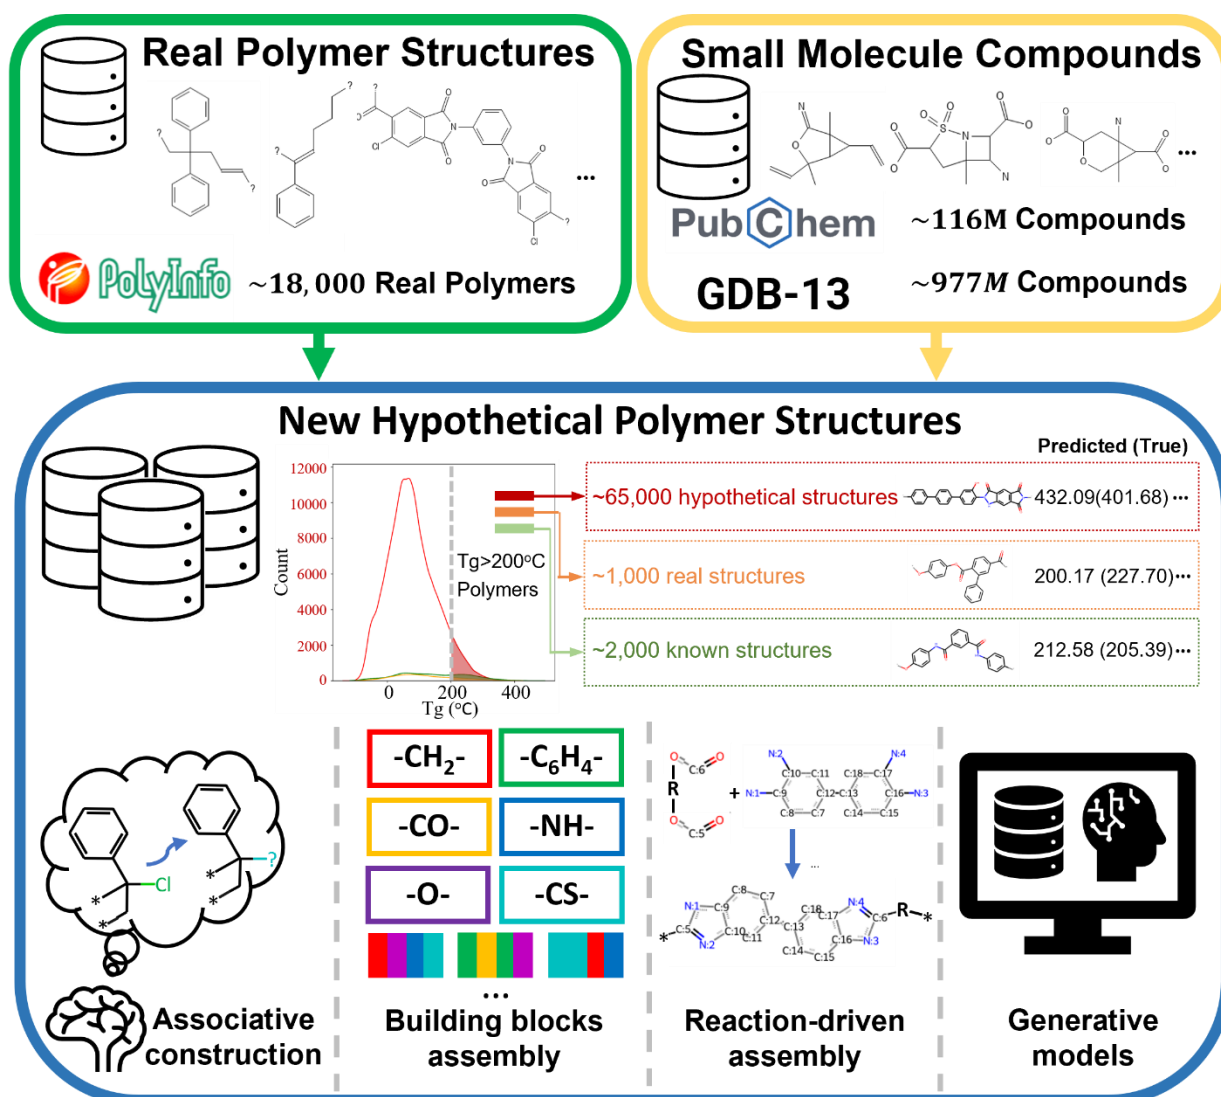


Figure 1. Hypothetical polymer structures play a crucial role in designing new polymers with exceptional properties. The four primary generation methods include manual design from existing structures, assembly of building blocks, utilizing existing small molecule compounds and synthetic routes, and the use of deep generative models.

To obtain candidates on a larger scale, another strategy for generating hypothetical polymer structures is based on existing small molecules and known polymerization reactions (or synthetic routes). As shown in **Figure 1**, the PolyInfo¹⁷ database lists merely 18,697 homopolymer structures. In comparison, there are around 116 million real small molecule compounds documented in PubChem¹⁸ and GDB-13¹⁹ offers us more than 900 million hypothetical small molecule compounds, which provide us with a vast chemical space for drug discovery. Taking advantage of these existing small molecules, Tao et al. generated an 8 million hypothetical polyimides and uncovered polyimides that possessed a multitude of outstanding thermal and mechanical properties simultaneously.^{8, 20} Using diamine and dianhydride monomers sourced

from PubChem, hypothetical polyimides were generated following a predefined reaction route. To efficiently screen these generated compounds, a machine learning method was employed for high-throughput evaluation. In a similar vein, Wang et al. generated 110 hypothetical polyimides by utilizing diamine and dianhydride monomers, resulting in high-temperature polymer dielectrics.⁶ This approach can provide a large number of candidates, but its chemical space is still limited by the small molecules used.²¹

With the rise of deep learning, generative models, and reinforcement learning, an increasing number of researchers are utilizing deep generative methods to expand the chemical space of various materials. This trend is particularly evident in the fields of cheminformatics and drug discovery.²²⁻⁴⁵ In polymer informatics, Ma and Luo created the PI1M dataset, comprising 1 million hypothetical polymers generated using an RNN trained on actual polymers sourced from PolyInfo.⁴⁶ In their study, they compiled 12,000 homopolymer structures from the PolyInfo database to train an RNN model. This training enabled the generation of 1 million new polymers, collectively referred to as PI1M. It was observed that while PI1M encompasses a chemical space similar to PolyInfo, it also fills in gaps where PolyInfo data is lacking, thereby offering a more comprehensive view of the polymer landscape. **Figure 1** illustrates how the ML-generated PI1M dataset provides researchers with a larger pool of promising hypothetical candidates for polymer informatics studies.⁴⁷

The other researchers directly generate hypothetical polymers with tailored properties using these deep generative models. For example, Wu et al. introduced Bayesian molecular design to discover polymers with high thermal conductivity.⁷ Gurnani et al. employed graph-to-graph (G2G) translation, called polyG2G, which can discern subtle chemical differences (referred to as translations) leading to significant property variations in polymeric materials.⁴⁸ A latent space searching strategy is employed in this study to generate hypothetical polymers with desired properties. They then used this knowledge to sample and design new polymers with high E_g and electron injection barrier. Batra et al. utilized syntax-directed VAE in conjunction with Gaussian process regression (GPR) models to identify polymers expected to exhibit robustness under extreme conditions, such as high temperatures, high electric fields, and their combination.⁴⁹ Liu et al. employed an invertible graph generative model to generate hypothetical polymers with promising properties, particularly focusing on high-temperature polymer dielectrics.⁵⁰ Kim et al. employed a method of searching and decoding within the latent space offered by a VAE to generate candidates with high polymer Log P^{51} values.⁵² Liu et al. developed a surrogate deep neural network model to predict thermal conductivity and compiled a library of polymer units consisting of 32 sequences. They utilized two advanced multi-objective optimization algorithms: Unified Non-dominated Sorting Genetic Algorithm III and Q-Noisy Expected Hypervolume Improvement, for designing sequence-ordered polymers that not only exhibit high thermal conductivity but also possess feasible synthetic potential.⁵³

When researchers intend to employ generative models in *de novo* polymer design, the initial step involves selecting a suitable model. However, at present, there is no work dedicated to assisting in the selection of generative models for hypothetical polymer structures. In contrast, numerous studies have been conducted to compare the performance of various models on small drug-like molecules, greatly aiding researchers in the field of drug discovery. One notable example of such a benchmarking platform is Molecular Sets (MOSES), which was developed to standardize the training and comparison of generative models for small molecules.²⁴ Zhang et al. conducted a benchmark study with a focus on functional groups and ring systems.⁵³ Weng et al. performed a benchmark specifically centered around biological properties.⁵⁴ Recently, Nigam et al. created a set of practical benchmark tasks called "Tartarus", which relies on physical simulations of molecular systems to emulate real-world challenges in molecular design for materials, drugs, and chemical reactions.⁵⁵

However, the complexity and unique properties of polymers necessitate distinct approaches and considerations compared to small molecules. This means that the conclusions derived from tests based on small molecules cannot be directly applied to the generation of hypothetical polymer structures. Therefore, there is an urgent need for developing specific benchmarks and methodologies tailored to the unique challenges and requirements in the molecular design of polymers. In this study, we initially used three different polymer datasets: real polymers from PolyInfo¹⁷, and hypothetical polyimides generated based on GDB-13¹⁹ and PubChem^{56, 57}, to train six generative models - VAE, AAE, ORGAN, CharRNN, REINVENT, and GraphINVENT. These models were trained on each dataset and generated about 10 million hypothetical polymer structures. We then evaluated these hypothetical polymer structures using the fraction of valid polymer structures f_v , the fraction of unique polymer structures from a sample of 10,000 f_{10k} , the Nearest Neighbor Similarity (SNN), the Internal Diversity (IntDiv) metric, and the Fréchet ChemNet Distance (FCD). These five metrics are provided by the MOSES platform. Furthermore, the t-Distributed Stochastic Neighbor Embedding (t-SNE) method are employed to visualize their chemical space distribution.

We further used reinforcement learning techniques, targeting the Tg, to train CharRNN, REINVENT, and GraphINVENT models to design hypothetical polymer structures with high Tg values. These three models are selected because of their outstanding performance based on the previous evaluation. All these models have demonstrated success in generating hypothetical polymers with high Tg values after 1,000-generation training. Overall, CharRNN provides us with the most favorable results. On the other hand, REINVENT stands out in terms of probability distribution but exhibits the lowest efficiency. The results of this study demonstrate the immense potential of generative models in the field of polymer informatics. They also provide valuable insights into the capabilities and limitations of various generative models within the realm of polymer science and

engineering. This understanding is crucial for researchers when it comes to selecting the most appropriate generative model for their specific needs.

2. RESULTS And DISCUSSION

2.1 DATASET AND DEEP GENERATIVE MODELS

Three datasets were employed, including real homopolymers manually collected from PolyInfo, and hypothetical polyimides generated using small molecules (polycondensation between diamine and dianhydride/diisocyanate monomers) from PubChem and GDB-13, as discussed in our previous study.²⁰ The real polymer dataset contains about 13,000 structures, while generative models typically require more training data. For example, Polykovskiy et al. utilized approximately 4.5 million samples for their work on MOSES²⁴, and Zhang et al. used around one million samples for their study⁵³. Therefore, we also utilized two hypothetical polyimides datasets for this purpose. The hypothetical polyimides generated using small molecules from PubChem and GDB-13 include a large number of structures, and we randomly selected around 10 million for model training.

Besides the difference in the number of samples in these datasets, these three datasets also vary in molecular weight and the number of types of atoms. It's important to note that the molecular weight values mentioned refer specifically to the molecular weight of the repeat units, not the entire polymer. This distinction is crucial because the repeat units serve as the input for analysis and modeling in these studies. The molecular weight of the entire polymer would be significantly higher and varies depending on the number of repeat units in the polymer chain.

Repeat units of real polymers from the PolyInfo database exhibit an average molecular weight of 443.7 and an average of 34.1 atoms per sample, encompassing 25 different types of atoms. In contrast, repeat units of hypothetical polyimides derived from PubChem show an average molecular weight of 530.4 and an average of 40.7 atoms, but with a limited variety of only 5 types of atoms. Repeat units of hypothetical polyimides created based on the GDB-13 have a higher average molecular weight of 645.8 and an average of 48.5 atoms per sample, featuring 18 different types of atoms. **Table S1 to S3** in the Supporting Information provide a detailed count of each atom type present in these datasets. These factors could significantly impact the training and performance of generative models. Specifically, the average number of atoms directly affects the size of the strings and graphs used for network input, while the variety of atomic types influence the molecular design of polymers by using these deep generative models. Utilizing these three diverse datasets enables us to better explore how different generative models perform in polymer informatics.

Polymer-Simplified Molecular Input Line Entry System (p-SMILES) strings are specialized string representations used to depict the chemical structures of polymers. These strings are instrumental in data-driven tasks related to polymer discovery, design, or prediction. The format

of a p-SMILES string is based on the standard SMILES syntax as defined by OpenSMILES⁵⁸. However, p-SMILES introduces a unique feature to represent polymers: it includes two stars ([*] or *) within the string. These stars signify the two endpoints of the polymer's repeat unit, effectively marking the bounds of the repeating segment in the polymer chain.

At present, large-scale generative models like Generative Pre-Trained Transformers (GPT)⁵⁹ have attracted widespread attention, but their scale and cost may be daunting for some researchers, particularly those who only wish to obtain some candidates in polymer design research. In these cases, smaller-scale generative models are still a more practical and accessible option. At the same time, due to the inherent differences between polymers and small molecules, such as higher complexity, larger molecular weight, and the use of p-SMILES, not all techniques applicable for generative models of small molecules are suitable for the generation of polymer structures. In this study, as shown in **Figure 2**, we selected the following six networks: VAE, AAE, ORGAN, CharRNN, REINVENT, and GraphINVENT, which are discussed in the following part.

VAE. VAE is a class of machine learning models that focuses on data generation and latent space learning. As shown in **Figure 2. (a)**, a VAE consists of two main components: the encoder and the decoder. The input encoder takes data x and maps it to a latent space representation, characterized by a distribution with mean μ and variance σ^2 . The VAE imposes a regularization by encouraging the latent distribution to resemble a standard Gaussian distribution $N(0, I)$, where I is the identity matrix. This is expressed in the objective function as the maximization of the similarity $\max \text{sim}(N(\mu, \sigma^2), N(0, I))$, which typically involves minimizing the Kullback-Leibler divergence between the two distributions. From the latent space, a sample z is drawn and passed to the decoder, which attempts to reconstruct the original input, producing \hat{x} . The training process involves minimizing the reconstruction error $\min \|\hat{x} - x\|$, making the decoded output as close as possible to the original input data. Our VAE model is implemented using the MOSES package.

AAE. AAE is a machine learning model that merges the concepts of Autoencoders (AE) and Generative Adversarial Networks (GANs). It can be observed from **Figure 2. (b)** that in the AAE framework, the input encoder receives raw data x and encodes it into a latent representation z . This latent representation is intended to follow a predefined probability distribution, typically a Gaussian distribution characterized by mean μ and variance σ . The output z from the encoder is then passed to the decoder, whose task is to reconstruct the input x to produce \hat{x} , with the goal of minimizing the reconstruction error $\min \|\hat{x} - x\|$. During training, the reconstructed output incrementally approaches the original input. Concurrently, the AAE includes a discriminator, which distinguishes whether the latent representations z generated by the encoder follow the set distribution. We implement the AAE model using the MOSES package as well.

ORGAN. ORGAN is a variant of the traditional GAN that incorporates objective reinforcement for improved generation of complex data. In the ORGAN framework, the generator creates synthetic data (represented by \mathbf{z}) which is intended to mimic real data samples. The discriminator, on the other hand, evaluates the synthetic data against real samples. Its goal is to distinguish between the two, effectively learning to tell apart genuine data from the imitations created by the generator. The twist in ORGAN compared to a standard GAN is the inclusion of a reinforcement signal, denoted by λ , which adjusts the generator's objectives beyond merely fooling the discriminator. The ORGAN is implemented with the MOSES package.

CharRNN. CharRNN is a type of neural network specifically designed for sequence prediction problems. This architecture is particularly useful for handling tasks where the input and/or output is a sequence of characters, such as text generation or in this case, p-SMILES generation. As shown in **Figure 2. (e)**, a CharRNN utilizes either Long Short-Term Memory (LSTM) or Gated Recurrent Unit (GRU) cells, both of which are variants of RNNs that are capable of learning long-term dependencies. The network depicted here processes the input sequence one character at a time (x_1, x_2, x_3, \dots), with each character being fed into the LSTM/GRU cell. These cells then produce an output sequence, where each output character is influenced by the previous characters in the sequence of p-SMILES. We use the MOSES package to implement the CharRNN.

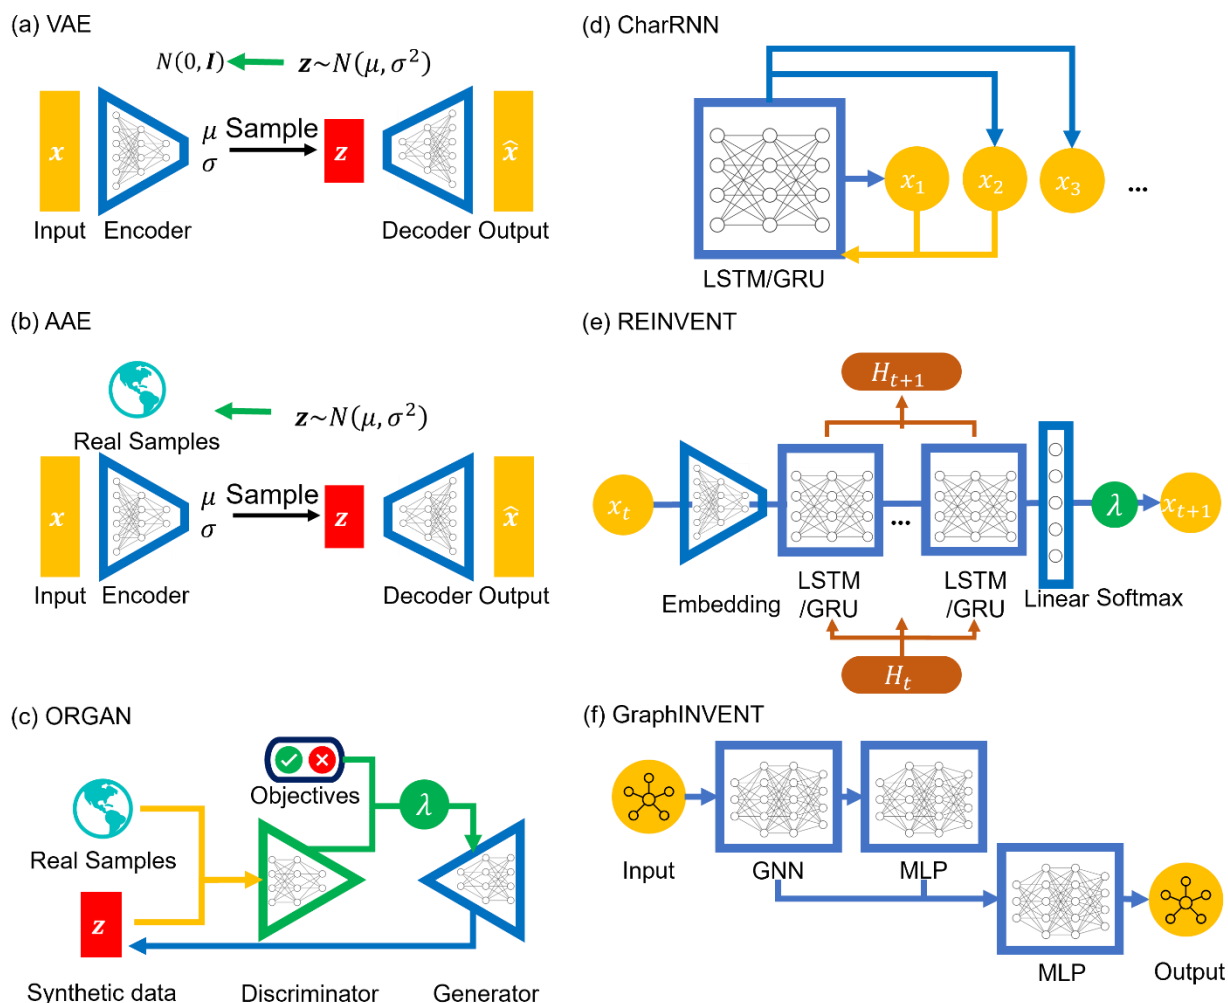


Figure 2. Architectures of six types of deep generative models: (a) VAE, (b) AAE, (c) ORGAN, (d) CharRNN, (e) REINVENT, and (f) GraphINVENT.

REINVENT. REINVENT is a sequence-based generative model that utilizes reinforcement learning for the generation of novel chemical entities, such as drug molecules or polymers. At the beginning of the sequence generation, an input x_t is processed through an input embedding layer, which transforms the discrete chemical symbols into continuous vectors. These vectors are then fed into a series of LSTM or GRU layers. Both LSTM and GRU are types of RNN cells capable of capturing long-term dependencies in sequential data. The recurrent cells process the input sequence, maintaining an internal state H_t that contains information about the sequence processed thus far. This state is updated with each new input symbol and is used to predict the next symbol in the sequence. After the LSTM/GRU layers, a linear layer followed by a softmax activation function produces a probability distribution over possible next symbols λ , from which the next symbol x_{t+1} is sampled.²⁶

GraphINVENT. GraphINVENT is a neural network model designed for generating novel molecular structures represented as graphs. The process starts with the input, which in the context of polymer chemistry, could be a set of initial monomer units or existing polymer fragments. The first stage of the model employs a Graph Neural Network (GNN), which processes the molecular graph as an input. After the GNN processes the molecular graph, the information is passed to a Multilayer Perceptron (MLP), which serves to interpret the features extracted by the GNN and assists in the decision-making process for the subsequent steps in molecule generation. The model includes an additional MLP layer that works in parallel with the first MLP to process separate features or to impose additional constraints or objectives on the generated structures.^{60, 61}

2.2 EVALUATION METRICS OF DEEP GENERATIVE MODELS

Once each model was trained, 10 million hypothetical polymers were sampled from each trained model. To demonstrate their effectiveness, we carefully selected five crucial metrics from the MOSES platform to evaluate the performance of these generative models. These metrics include the fraction of valid polymer structures f_v , which measures the percentage of chemically valid structures generated by the model. The chemical validity of the hypothetical polymer structure is determined by using RDKit package and the count of *, as [*] is exactly two for p-SMILES of homopolymers. We also considered the fraction of unique polymer structures from a sample of 10,000 f_{10k} , assessing the model's ability to generate diverse chemical structures.

Additionally, the Nearest Neighbor Similarity (SNN) was used to calculate the average similarity of the generated polymers to the closest polymer in the test set, providing an insight into how the generated polymers compared to known structures. SNN represents the average Tanimoto similarity $T(m_G, m_R)$. This similarity is calculated between the fingerprints of a polymer m_G in the generated set G and its closest neighboring polymer m_R in the reference dataset R :

$$\text{SNN}(G, R) = \frac{1}{|G|} \sum_{m_G \in G} \max_{m_R \in R} T(m_G, m_R),$$
$$T(m_G, m_R) = \frac{n_{m_G \& m_R}}{n_{m_G} + n_{m_R} + n_{m_G \& m_R}}$$

where n_X is the count of bits “on” in polymer X 's fingerprint but not in polymer Y 's fingerprint, and $n_{X \& Y}$ is the count of bits “on” both in polymer X 's fingerprint and in polymer Y 's fingerprint.

The Internal Diversity (IntDiv) metric, representing the average pairwise similarity among generated polymers, was included to gauge the diversity within the generated polymer structures.⁶² IntDiv assesses the chemical diversity within the generated set of polymers G :

$$\text{IntDiv}_p(G) = 1 - \sqrt[p]{\frac{1}{|G|^2} \sum_{m_1, m_2 \in G} T(m_G, m_R)^p}$$

Lastly, the Fréchet ChemNet Distance (FCD) was employed to quantify the difference in the distribution of the last layer activations of ChemNet⁶³, which is trained to predict bioactivities of about 6000 assays available in three major drug discovery databases (ChEMBL⁶⁴, ZINC⁶⁵, PubChem⁵⁶), effectively measuring the disparity between the generated polymer distribution and a reference set.⁶⁶ For two sets of polymers G and R , FCD is defined as

$$\text{FCD}(G, R) = \|\mu_G - \mu_R\|^2 + \text{Tr} \left[\Sigma_G + \Sigma_R - 2(\Sigma_G \Sigma_R)^{\frac{1}{2}} \right]$$

where μ_G , μ_R are mean vectors and Σ_G , Σ_R are full covariance matrices of activations for polymers from sets G and R , respectively.

After all, both t-SNE and Tanimoto similarity metrics were employed to assist in comparing differences between various polymer structures. T-SNE, a widely used technique for nonlinear dimensionality reduction and data visualization, effectively maintains nonlinear similarities between data points. It operates by initially determining the similarity between high-dimensional data points using a Gaussian distribution. Subsequently, it assesses the similarity among data points in a reduced, low-dimensional space based on a t-distribution. The goal of t-SNE is to minimize the disparity between these high-dimensional and low-dimensional similarities. These selected metrics will be employed in the initial phase of comparison to sieve out the generative models that demonstrate superior performance.

2.3 PERFORMANCE AND COVERAGE OF GENERATIVE MODELS

Figure 3 shows the performance of the six generative models when applied to the real polymer dataset from PolyInfo. In terms of f_v , the CharRNN model achieved the highest result, nearly reaching 0.9. Both the GraphINVENT and REINVENT models achieved greater than 0.5. However, the VAE, AAE, and ORGAN models obtained notably lower scores. These outcomes indicate a comparatively lower effectiveness of these models in generating valid polymer structures (or p-SMILES) compared to other models. For the metric f_{10k} , AAE, ORGAN, REINVENT, and VAE exhibit good performance, with scores around 0.8. CharRNN and GraphINVENT, while not performing as well as the aforementioned models, still achieve results greater than 0.5, which is considered acceptable.

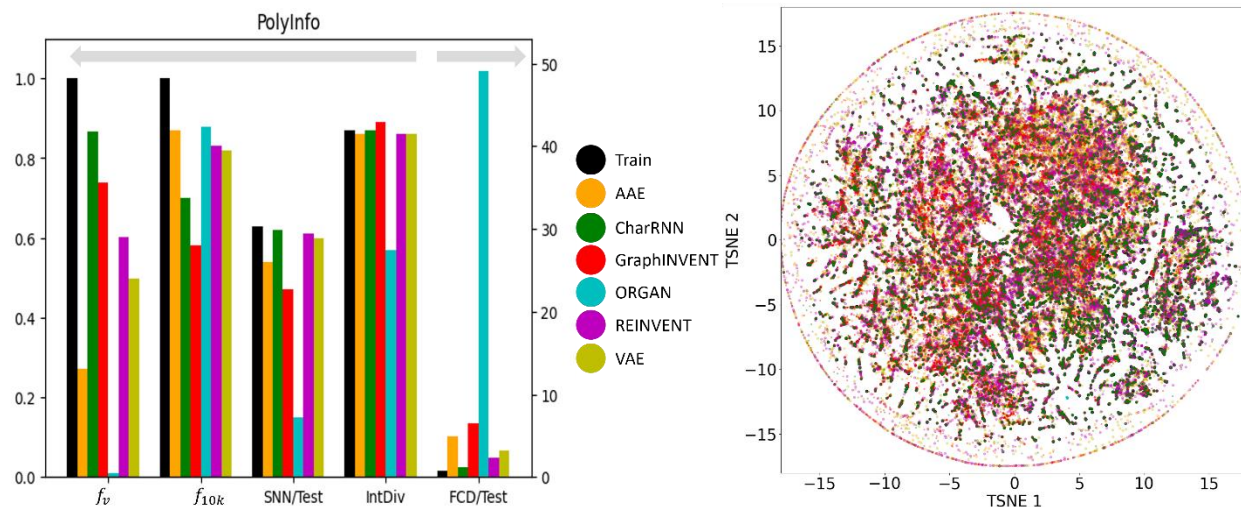


Figure 3. A comparison of the performance of the six generative models on the real homopolymer dataset collected from the PolyInfo, as well as the chemical space distribution of the generated polymers.

In evaluating the performance of generative models using the SNN and IntDiv metrics, higher values are generally sought after. These metrics provide insights into the models' ability to generate both diverse and chemically relevant polymer structures. It can be observed that all models, except for ORGAN, exhibit results that closely resemble those in the training set.

For the FCD metric, lower values are generally preferred. This metric measures the difference in distributions between the generated polymers and a reference set, with a lower score indicating that the generated polymers are more chemically similar to real polymers. The observations indicate that, similar to the SNN metric, VAE, REINVENT, and CharRNN achieved relatively low FCD scores. AAE and GraphINVENT obtained higher scores, while ORGAN exhibited a significantly higher FCD score.

Considering all the metrics collectively, it appears that CharRNN, REINVENT, and GraphINVENT deliver the best performance, while AAE and VAE follow closely behind. However, ORGAN's performance leaves much to be desired. This result bears similarity to previous benchmark work based on small molecules. In MOSES, Polykovskiy et al. found that among a wide array of models, CharRNN currently outperforms others in terms of these key metrics.²⁴ In RediscMol, Weng, G. et al. observed that CharRNN, VAE, and REINVENT yield superior results, followed by AAE and ORGAN.⁵⁴ Additionally, in studies considering ring system coverage and functional group coverage, AAE, REINVENT, VAE, CharRNN, and GraphINVENT all exhibit better performance compared to ORGAN.⁵³

CharRNN consistently shows remarkable results in these benchmark studies, while the performance of AAE and VAE tends to be less impressive in our result. This could be attributed to

the fact that the PolyInfo dataset is significantly smaller than datasets for small molecules. Additionally, the structural differences between real polymers and small molecules also play a role. The ZINC Clean Leads⁶⁷ used in the MOSES project have molecular weights ranging from 250 to 350 Daltons.²⁴ However, the molecular weight in the real polymer dataset varies widely, ranging from 14 to 2202 Daltons. This variation is due to the presence of polymers with complex structures as well as those with very simple repeat units. For example, polyethylene, the simplest polymer, has a p-SMILES representation of just '*C*'. The t-SNE visualization further corroborates the analysis derived from these metrics, providing a graphical representation of how well each model captures the chemical space of polymers. The individual t-SNE results for each model can be found in the Supporting Information.

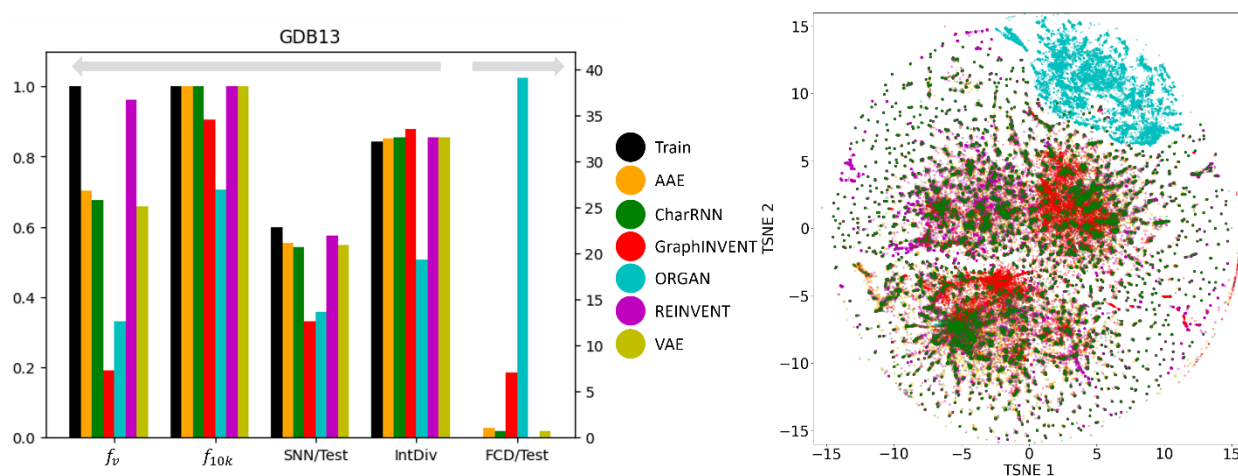


Figure 4. A comparison of the performance of the six generative models on the hypothetical polyimide dataset synthesized based on GDB-13, as well as the chemical space distribution of the generated polymers.

Figure 4 presents the performance of six different generative networks when applied to the hypothetical polyimide dataset based on GDB-13. For f_v , the REINVENT model achieved the highest result, nearly equal to 1. In comparison, the AAE, VAE, and CharRNN models show a similar performance level, with their values clustered around 0.7. On the other hand, the ORGAN and GraphINVENT models have considerably lower scores, below 0.2.

In f_{10k} part, several models exhibited impressive results. REINVENT, AAE, VAE, and CharRNN all achieved a result of 1. It indicates an excellent ability of these models to generate a diverse set of polymer structures, with no duplicates in a sample of 10,000 p-SMILES strings. GraphINVENT, while not reaching a perfect score, still performed commendably, with its value being close to 0.9. However, ORGAN scored below 0.7, indicating less diversity in its generated polymer structures.

For the SNN metric, it was observed that, apart from ORGAN and GraphINVENT, the other four models showed similar performance. Regarding IntDiv, all models except for ORGAN exhibited

closely matched performances. These observations suggest that REINVENT, AAE, VAE, CharRNN, and GraphINVENT are capable of producing a wide variety of polymer structures, demonstrating a good internal diversity among the generated hypothetical polymers.

Observations show that, similar to the SNN metric, the models AAE, VAE, REINVENT, and CharRNN achieved relatively low FCD scores. GraphINVENT recorded a somewhat higher FCD score, indicating less chemical similarity between its generated structures and the training dataset. ORGAN exhibited a significantly higher FCD score, implying a larger disparity between its generated structures and the real-world polymers.

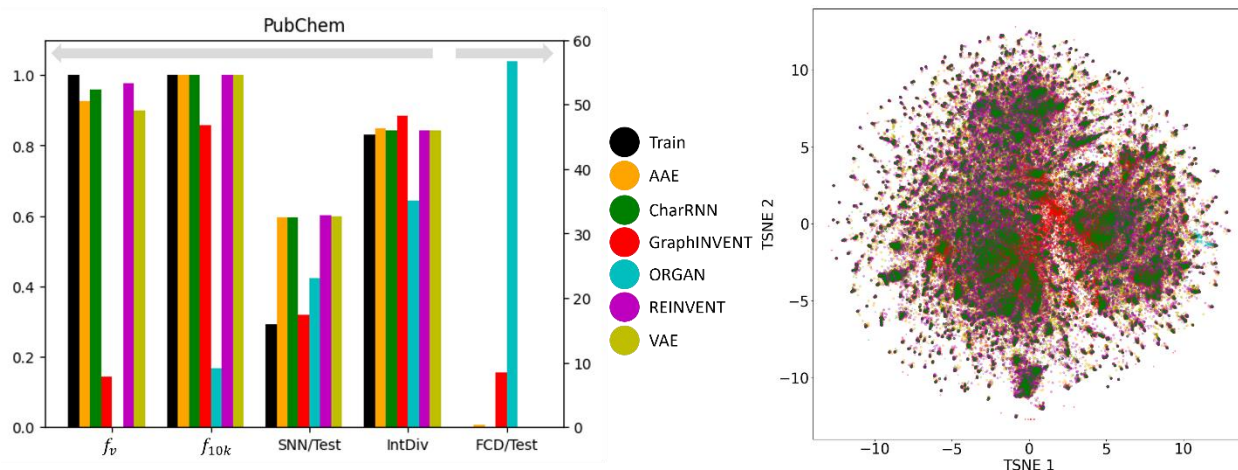


Figure 5. A comparison of the performance of the six generative models on the hypothetical polyimide dataset synthesized based on PubChem, as well as the chemical space distribution of the generated polymers.

Figure 5 denotes the performance of the same six generative models when applied to the hypothetical polyimide dataset derived from PubChem. It is observed that the performance and comparative results of these six models are almost consistent with those outcomes from training on the hypothetical polyimide based on GDB-13, which means the REINVENT model demonstrated the best performance. However, the performance of the ORGAN model was notably worse, to the point of being considered unacceptable for the task at hand.

All the above results show that the REINVENT model shows the most favorable performance. The AAE, CharRNN, and VAE models follow closely, while GraphINVENT and ORGAN demonstrate a much worse performance. It should be noted that the performance of these six generative models vary across the three datasets, which is related to the characteristics of each dataset. Compared to the real polymer dataset, PolyInfo, the AAE and VAE show significant improvement on the hypothetical polyimide datasets. This suggests that increasing the amount of training data helps improve the performance of AAE and VAE. Moreover, these generative models perform better on

the dataset based on PubChem than the one based on GDB-13, indicating their suitability for scenarios with fewer types of atoms.

GraphINVENT performs better on the real polymer dataset compared to the other two datasets. This is likely due to the real polymer dataset having a smaller average molecular weight and number of atoms, resulting in smaller and simpler graph data structures. As for the two RNN-based generative networks, REINVENT demonstrates outstanding performance on both hypothetical polyimide datasets, proving its ability to handle such tasks when sufficient data is available. CharRNN, however, shows weaker performance on the polyimide dataset based on GDB-13 compared to the other two datasets. This is attributed to the largest average molecular weight and number of atoms in the polyimide dataset devised from GDB-13. RNN networks process input and output one unit at a time using GRU or LSTM, and the task becomes more challenging as the string length increases.

We also observed that these findings align closely with the results of work of Zhang et al., particularly in their results regarding ring system coverage. In their study, they utilized the GDB-13 dataset as a training set, which happens to be one of the sources we used to generate hypothetical polyimides for our research.⁵³

2.4 DEEP GENERATIVE DESIGN WITH REINFORCEMENT LEARNING

The introduction of reinforcement learning algorithms empowers generative models with the capability to design hypothetical polymer structures possessing specific properties. This method represents a transformative step towards more efficient and purpose-driven material discovery and design. **Figure 6** illustrates the fundamental architecture of reinforcement learning as applied to these generative models. In this framework, the agent, which is the generative model, initiates the process by generating a set of candidate polymer structures. The evaluation of these candidates follows a specific scoring mechanism.

Firstly, the generated p-SMILES strings are converted into molecular fingerprints (MF). These MFs are then used as input to a Feed-Forward Neural Network (FNN), which is tasked with predicting the T_g values of these candidates. Detailed information about the FNN is available in the Supporting Information. After obtaining the T_g predictions, a sigmoid function is applied to these values. The output of this sigmoid function is treated as the reward, which is fed back to the agent. The feedback received in the form of rewards is then used by the agent to further train and optimize its performance. According to the previous results, particularly the comparison of generative models trained on the PolyInfo dataset, we selected REINVENT, CharRNN, and GraphINVENT as our models of choice for generative design of new polymers. Employing reinforcement learning, we used PolyInfo as the training dataset with the goal of training these models to generate hypothetical polymer structures that exhibit high T_g values.




Figure 6. The core framework of reinforcement learning with deep generative model, and the specific data flow utilized in this study.

Figure 7 presents the performance of these three models undergoing reinforcement learning. The leftmost part of the figure shows the change in the predicted average T_g of the generated hypothetical polymer structures across training generations, along with the predicted T_g distribution of polymers generated at the 200th, 600th, and 1000th training steps as well as the training set. As training iterations increased, it was observed that the predicted T_g values of the hypothetical polymer structures generated by all three generative models showed an upward trend. Notably, CharRNN achieved the highest average predicted T_g value at the 1000th step, while REINVENT and GraphINVENT exhibited similar performance. Additionally, the distribution of the predicted T_g values for all generated hypothetical polymers shifted towards higher values. This outcome demonstrates the capability of reinforcement learning to effectively steer the generative process towards specific target properties, in this case, achieving higher T_g in the hypothetical polymer structures.

In the middle of **Figure 7**, t-SNE plot, the chemical space covered by the training set is represented in grey, while the transition from deep purple to yellow indicates the chemical space of polymer structures generated over training epochs ranging from 0 to 1200. This color gradient visually represents the evolution of the generated polymers' chemical space throughout the reinforcement learning process.

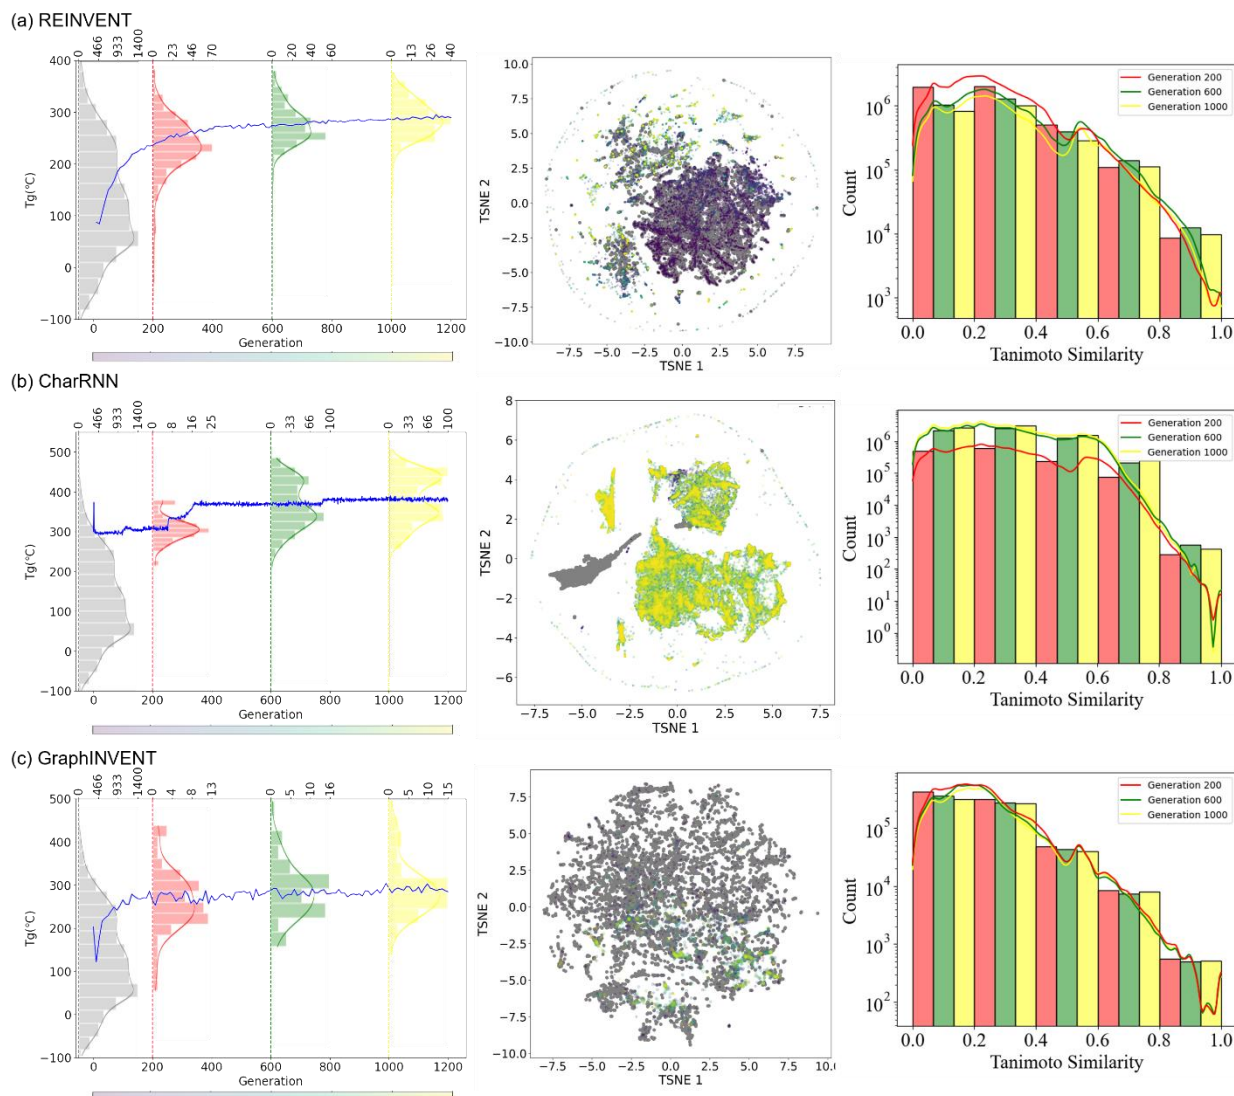


Figure 7. Performance of three generative models, REINVENT, CharRNN, and GraphINVENT, combined with reinforcement learning. (Left) There's a depiction of how the predicted average T_g of the generated hypothetical polymer structures evolves over various training generations. This includes a detailed view of the T_g distributions for structures produced at the 200th, 600th, and 1000th training steps, as well as those in the training set. (Middle) The central t-SNE plot visualizes the chemical space: the area covered by the training set is shown in grey, while the progression from deep purple to yellow represents the evolving chemical space of the polymer structures generated throughout the training epochs, from the initial to the 1200th. (Right) The graphs displaying the Tanimoto similarity between the hypothetical polymer structures generated at the 200th, 600th, and 1000th steps and the training set provides critical insights into the dynamics of the training process.

For REINVENT and GraphINVENT, it was observed that the chemical space of the newly generated polymers remained within the bounds of the chemical space covered by the original training set. As training progressed, there was a noticeable shift from larger purple-red regions to smaller, more concentrated yellow areas. Similar to REINVENT and GraphINVENT, the CharRNN model also exhibited a gradual concentration of the chemical space of the generated polymer structures during the training process. However, a distinct behavior was observed in CharRNN's approach. Unlike the other two generative models, CharRNN began within the chemical space covered by the original training set and progressively expanded its search into chemical spaces beyond what was covered in the training set. As a result, the hypothetical polymers generated by CharRNN occupied a much larger area in the chemical space.

The right panel of **Figure 7** illustrates the Tanimoto similarity between the hypothetical polymer structures generated at the 200th, 600th, and 1000th training steps and the training set, reveals an important aspect of the training process. This observation suggests that, as the models are trained, the generated polymer structures maintain a certain level of structural resemblance to those found in the initial training set. The absence of a convergence towards zero in the Tanimoto similarity indicates that the models are not diverging significantly from the structural characteristics of real polymers.

This pattern suggests that as the number of training epochs increased, both REINVENT and GraphINVENT models started to focus on generating polymer structures within specific, more defined regions of the chemical space. This convergence towards certain areas within the training set's chemical space could indicate that the models are focusing on regions that are more likely to yield polymers with the desired high T_g values. This demonstrates that reinforcement learning strategies are effectively guiding generative models in exploring the polymer chemical space.

Meanwhile, it was observed that the results from the REINVENT and GraphINVENT models remained within the chemical space defined by the training dataset, while CharRNN showed an expansion beyond the initial training set boundaries. It's important to note that the different chemical space distributions observed do not affect the similarity of the generated hypothetical polymer structures to the training set. This is because the hypothetical polymer structures generated by the CharRNN, REINVENT, and GraphINVENT models exhibit a Tanimoto similarity that is essentially consistent with each other. The REINVENT and GraphINVENT models are particularly adept at controlling the generated structures within the confines of the training set, making them suitable choices for researchers who desire such candidates. As for CharRNN, the expansion beyond the initial training set boundaries suggests that it was exploring more novel regions of the chemical space, potentially leading to the discovery of new polymer structures with higher T_g values. This exploration outside the known chemical space is a key factor in why

CharRNN's generated hypothetical polymers had overall higher mean predicted T_g values during the training.

However, as previously discussed, the limited number of polymer structures in the real polymer dataset can lead to decreased effectiveness. Ma et al. utilized RNNs and a reinforcement learning algorithm to generate hypothetical polymer structures with high thermal conductivity. They used a significantly larger training dataset (PI1M), consisting of 1 million samples, which far exceeds the size of the real polymer dataset.⁴⁶ A larger dataset provides more comprehensive coverage of the chemical space in their study, allowing models to learn a wider range of patterns and features. This can lead to the generation of more unique and diverse polymer structures, enhancing the potential for discovering novel materials with desirable properties. In their study, the visualization of generated polymers alongside real polymers from PolyInfo using t-SNE showed a pattern similar to that of CharRNN's results here. The newly generated molecules show an expansion beyond the initial boundaries of the training set, indicating exploration into new areas of the chemical space. Additionally, Moret et al. have demonstrated the generation of novel small molecules with bespoke properties and structural diversity using an RNN. The chemical space explored in their research exhibits a pattern similar to what we have observed in other studies, highlighting the RNN's capability in navigating and innovating within the chemical space.

Furthermore, graph-based generative networks for polymers, such as PolyG2G, have also exhibited outstanding performance.⁴⁸ Instead of using reinforcement learning, their network employed a latent space searching strategy to generate hypothetical polymers with desired properties. The same concept of latent space utilization is also evident in VAEs based on the inclusion of a latent space in these models' architecture.^{49, 52} Similarly, Liu et al. utilized a graph-based invertible molecular generative model along with a latent space strategy for the design of high-temperature polymer dielectrics.⁵⁰ Observations of these two graph-based generated models employing latent space strategies reveal that the main frameworks of the generated repeat units often bear resemblance to certain structures within the training set. This similarity might be a contributing factor to the close alignment of the chemical space of the generated hypothetical polymers with that of the training set.

The size of the training dataset also significantly impacts the efficiency and uniqueness of the hypothetical polymer structures generated by these models, both of which are critical factors for practical applications. **Figure 8** illustrates the comparison of these two metrics – the efficiency and non-redundancy rates – for the three networks at the 200th, 600th, and 1000th steps, as well as their overall trends throughout the training process.




Figure 8. Effect of iteration steps on the validity and redundancy rates of three generative models, CharRNN, GraphINVENT, and REINVENT, during the reinforcement learning process.

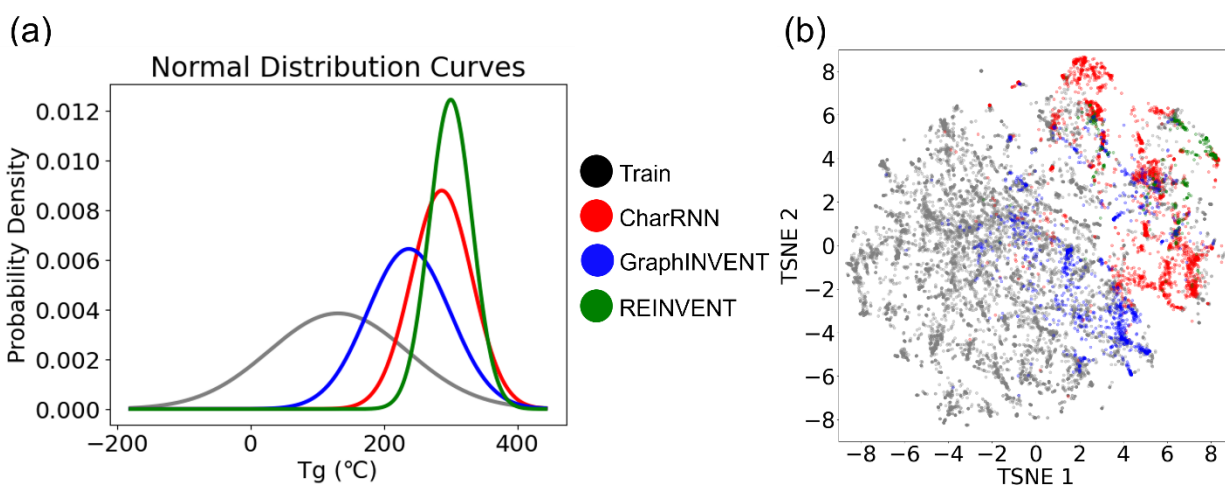


Figure 9. (a) Normalized probability density distribution of predicted T_g values and the (b) chemical space distribution of the hypothetical valid unique polymers generated by CharRNN (red), GraphINVENT (blue), REINVENT (green), and the real polymers.

Then the 1,000th-training-iteration models are used for 100,000 hypothetical polymers generation. **Figure 9 (a)** displays the normalized probability density distribution of predicted T_g values of these hypothetical valid unique polymers. When the models are employed to generate a large number of hypothetical polymer structures, there is a slight shift in the mean prediction values. It is evident that the probability distributions of the three generative models are significantly different from the training set, favoring higher T_g values. Generative models based on RNN and GNN architectures have been effectively used to directly create hypothetical polymer structures with desired properties, achieving commendable results. Among them, REINVENT has the highest mean and the smallest variance in its probability density distribution, indicating that it can more stably generate many high T_g hypothetical polymer structures. CharRNN and GraphINVENT are less effective in comparison. However, it is important to note, as shown in **Figure 8**, that the

unique rate and validity rate of REINVENT model are relatively low. In contrast, CharRNN is considered as the best option.

Figure 9 (b) showcases the chemical space distribution of these hypothetical, valid, and unique polymers. The results align with those presented in **Figure 7**, showing that among the three generative models, CharRNN generates hypothetical polymers (red points) that are the most distinct from real polymers (grey points) in terms of distance. In contrast, the polymers generated by the other two models are interspersed within the distribution of real polymers.

From these results, it is evident that CharRNN demonstrates a distinct advantage in both efficiency and uniqueness. This superiority is likely connected to its broader exploration of the polymer chemical space. As previously discussed in comparisons and discussions of various generative models, CharRNN has shown the best performance with real polymers collected from PolyInfo and a variety of small molecule tests. While REINVENT exhibits the best normalized probability density distribution, its unique rate should be considered. Hence, REINVENT becomes the optimal choice specifically when there is a requirement to generate a substantial volume of candidates, ranging from hundreds to thousands.

3 CONCLUDING REMARKS

This study conducts a comprehensive evaluation of generative models within the context of polymer informatics, highlighting both their potential and limitations. Initially, six generative models – AAE, VAE, CharRNN, REINVENT, GraphINVENT, and ORGAN – were tested and trained using datasets of hypothetical polyimides based on PubChem and GDB-13, as well as real polymer datasets collected from PolyInfo. The performance of these generative models was assessed using various metrics: the fraction of valid structures, the fraction of unique structures from a sample of 10,000, SNN, IntDiv, and FCD. It was observed that CharRNN, REINVENT, and GraphINVENT produced superior results when trained with the PolyInfo dataset. Meanwhile, REINVENT demonstrated outstanding performance when trained with the two hypothetical polyimide datasets, with AAE, VAE, and CharRNN also showing commendable outcomes. This difference in performance may be attributed to the more complex structures and larger molecular weights of hypothetical polyimides.

Subsequently, CharRNN, REINVENT, and GraphINVENT were combined with reinforcement learning algorithms using the PolyInfo dataset to generate hypothetical polymer structures with higher T_g values. All three models performed impressively, but with notable differences in their capabilities. CharRNN displayed a unique ability to extend beyond the chemical space of the training set, generating polymers with higher predicted T_g values. After training, REINVENT demonstrates the most outstanding probability distribution in its generated results. However, compared to CharRNN and GraphINVENT, it has a lower unique rate and valid outcomes.

The study underscores the need for specific benchmarks and methodologies tailored to the unique challenges of polymer design. The integration of reinforcement learning proved effective in guiding the generative process towards the desired properties, highlighting the potential of these models in future materials design and discovery. This work also leverages the power of computational modeling and machine learning, paving the way for more targeted and efficient development of new polymeric materials, such as organic photovoltaics, polymer membranes, and dielectrics.

4. EXPERIMENTAL PROCEDURES

4.1 MODEL TRAINING

For the PolyInfo dataset, approximately 11,000 homopolymers were randomly selected to constitute the training set, while around 1,200 were designated as the test set. Regarding the other two hypothetical polyimide datasets based on PubChem and GDB-13, a subset of 0.8 million polyimides was randomly chosen and utilized as the training set for all the generative models. Furthermore, two additional sets of 0.2 million polyimides were specifically selected to serve as the validation sets. For the implementation of REINVENT and GraphINVENT in this study, the hyperparameters were directly sourced from their respective GitHub repositories. In the case of CharRNN, AAE, VAE, and ORGAN, the hyperparameters were adopted from the models' configuration files available in the MOSES GitHub repository (<https://github.com/molecularsets/moses>).

The CharRNN model integrated with reinforcement learning was utilized, with its code accessible at the GitHub repository (<https://github.com/aspuru-guzik-group/Tartarus>). The REINVENT model, incorporating reinforcement learning, was employed, with its code available at <https://github.com/MolecularAI/Reinvent>. The GraphINVENT model, enhanced with reinforcement learning, was utilized in this study. Its code is accessible at <https://github.com/olsson-group/RL-GraphINVENT>. This ongoing process of generation, evaluation, and feedback allows the generative models to progressively improve in its ability to design hypothetical polymer structures that closely match the targeted properties, thus enhancing the efficiency and effectiveness of the materials design process.

4.2 TECHNICAL DETAILS

The training of generative models from the MOSES platform was conducted using the Docker container "molecularsets/moses". This training took place on Linux workstations equipped with NVIDIA Quadro RTX 8000 graphics cards, utilizing CUDA 12.1 for computational acceleration. For the GraphINVENT model, the training environment comprised Python 3.6.8 and PyTorch 1.3.1. This model was also trained on Linux workstations, but with NVIDIA Quadro P6000 graphics cards, again leveraging CUDA 12.1 for enhanced processing capabilities. Regarding the REINVENT model,

it was trained using Python 3.7.7 and PyTorch 1.7.0. This model's training was performed on Linux workstations equipped with NVIDIA RTX A6000 graphics cards, utilizing the same 12.1 version of CUDA, 12.1, for computational support.

5. SUPPLEMENTAL INFORMATION

Supplemental information can be found online at: XXX.

6. ACKNOWLEDGMENTS

We gratefully acknowledge financial support from the Air Force Office of Scientific Research through the Air Force's Young Investigator Research Program (FA9550-20-1-0183; Program Manager: Dr. Ming-Jen Pan and Capt Derek Barbee), Air Force Research Laboratory/UES Inc. (FA8650-20-S-5008, PICASSO program), and the National Science Foundation (CMMI-2314424, CMMI-2316200, and CAREER-2323108). Y.L. would also like to thank the support from 3M's Non-Tenured Faculty Award. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the U.S. Department of Defense or National Science Foundation. The authors also acknowledge the National Renewable Energy Laboratory for providing HPC resources that have contributed to the research results reported within this paper. Support for this research was also provided by the University of Wisconsin–Madison, Office of the Vice Chancellor for Research and Graduate Education with funding from the Wisconsin Alumni Research Foundation.

7. AUTHOR CONTRIBUTIONS

Conceptualization, Y.L. and V.V.; methodology, T.Y, L.T, and Y.L.; software T.Y, L.T.; validation, T.Y.; formal analysis, T.Y, Y.L.; investigation, T.Y, Y.L.; resources, Y.L.; data curation, T.Y, L.T.; writing—original draft, T.Y.; writing—review & editing, T.Y, V.V., and Y.L.; visualization, L.T.; supervision, Y.L.; funding acquisition, Y.L.

8. DECLARATION OF INTERESTS

The authors declare no competing interests.

REFERENCE

- (1) Li, G.; Zhu, R.; Yang, Y. Polymer solar cells. *Nature photonics* **2012**, *6* (3), 153-161.
- (2) Hsissou, R.; Seghiri, R.; Benzekri, Z.; Hilali, M.; Rafik, M.; Elharfi, A. Polymer composite materials: A comprehensive review. *Composite structures* **2021**, *262*, 113640.
- (3) Diahm, S. Polyimide in electronics: Applications and processability overview. *Polyimide for Electronic and Electrical Engineering Applications* **2021**, 2020-2021.
- (4) Gouzman, I.; Grossman, E.; Verker, R.; Atar, N.; Bolker, A.; Eliaz, N. Advances in polyimide-based materials for space applications. *Advanced Materials* **2019**, *31* (18), 1807738.

- (5) Anstey, A.; Chang, E.; Kim, E. S.; Rizvi, A.; Kakroodi, A. R.; Park, C. B.; Lee, P. C. Nanofibrillated polymer systems: Design, application, and current state of the art. *Progress in Polymer Science* **2021**, *113*, 101346.
- (6) Wang, R.; Zhu, Y.; Fu, J.; Yang, M.; Ran, Z.; Li, J.; Li, M.; Hu, J.; He, J.; Li, Q. Designing tailored combinations of structural units in polymer dielectrics for high-temperature capacitive energy storage. *Nature Communications* **2023**, *14* (1), 2406.
- (7) Wu, S.; Kondo, Y.; Kakimoto, M.-a.; Yang, B.; Yamada, H.; Kuwajima, I.; Lambard, G.; Hongo, K.; Xu, Y.; Shiomi, J.; et al. Machine-learning-assisted discovery of polymers with high thermal conductivity using a molecular design algorithm. *npj Computational Materials* **2019**, *5* (1), 66. DOI: 10.1038/s41524-019-0203-2.
- (8) Yue, T.; He, J.; Tao, L.; Li, Y. High-Throughput Screening and Prediction of High Modulus of Resilience Polymers Using Explainable Machine Learning. *Journal of Chemical Theory and Computation* **2023**.
- (9) Zhang, Z. P.; Rong, M. Z.; Zhang, M. Q. Polymer engineering based on reversible covalent chemistry: A promising innovative pathway towards new materials and new functionalities. *Progress in Polymer Science* **2018**, *80*, 39-93.
- (10) Chen, J.; Zhou, Y.; Huang, X.; Yu, C.; Han, D.; Wang, A.; Zhu, Y.; Shi, K.; Kang, Q.; Li, P. Ladderphane copolymers for high-temperature capacitive energy storage. *Nature* **2023**, *615* (7950), 62-66.
- (11) Dong, J.; Li, L.; Qiu, P.; Pan, Y.; Niu, Y.; Sun, L.; Pan, Z.; Liu, Y.; Tan, L.; Xu, X. Scalable Polyimide-Organosilicate Hybrid Films for High-Temperature Capacitive Energy Storage. *Advanced Materials* **2023**, *35* (20), 2211487.
- (12) Zhang, Y.; Zhang, J.; Suzuki, K.; Sumita, M.; Terayama, K.; Li, J.; Mao, Z.; Tsuda, K.; Suzuki, Y. Discovery of polymer electret material via de novo molecule generation and functional group enrichment analysis. *Applied Physics Letters* **2021**, *118* (22).
- (13) Yin, X.; Wan, T.; Deng, X.; Xie, Y.; Gao, C.; Zhong, C.; Xu, Z.; Pan, C.; Chen, G.; Wong, W.-Y. De novo design of polymers embedded with platinum acetylides towards n-type organic thermoelectrics. *Chemical Engineering Journal* **2021**, *405*, 126692.
- (14) Mei, D.; Yan, L.; Liu, X.; Zhao, L.; Wang, S.; Tian, H.; Ding, J.; Wang, L. De novo design of single white-emitting polymers based on one chromophore with multi-excited states. *Chemical Engineering Journal* **2022**, *446*, 137004.
- (15) Sharma, V.; Wang, C.; Lorenzini, R. G.; Ma, R.; Zhu, Q.; Sinkovits, D. W.; Pilania, G.; Oganov, A. R.; Kumar, S.; Sotzing, G. A. Rational design of all organic polymer dielectrics. *Nature communications* **2014**, *5* (1), 4845.
- (16) Li, H.; Chang, B. S.; Kim, H.; Xie, Z.; Lainé, A.; Ma, L.; Xu, T.; Yang, C.; Kwon, J.; Shelton, S. W. High-performing polysulfate dielectrics for electrostatic energy storage under harsh conditions. *Joule* **2023**, *7* (1), 95-111.
- (17) Otsuka, S.; Kuwajima, I.; Hosoya, J.; Xu, Y.; Yamazaki, M. PoLyInfo: Polymer database for polymeric materials design. In *2011 International Conference on Emerging Intelligent Data and Web Technologies*, 2011; IEEE: pp 22-29.
- (18) Kim, S.; Chen, J.; Cheng, T.; Gindulyte, A.; He, J.; He, S.; Li, Q.; Shoemaker, B. A.; Thiessen, P. A.; Yu, B. PubChem 2023 update. *Nucleic acids research* **2023**, *51* (D1), D1373-D1380.

- (19) Blum, L. C.; Raymond, J.-L. 970 million druglike small molecules for virtual screening in the chemical universe database GDB-13. *Journal of the American Chemical Society* **2009**, *131* (25), 8732-8733.
- (20) Tao, L.; He, J.; Munyaneza, N. E.; Varshney, V.; Chen, W.; Liu, G.; Li, Y. Discovery of multi-functional polyimides through high-throughput screening using explainable machine learning. *Chemical Engineering Journal* **2023**, *465*, 142949.
- (21) Ohno, M.; Hayashi, Y.; Zhang, Q.; Kaneko, Y.; Yoshida, R. SMiPoly: Generation of Synthesizable Polymer Virtual Library using Rule-based Polymerization Reactions. **2023**.
- (22) Bjerrum, E. J.; Threlfall, R. Molecular generation with recurrent neural networks (RNNs). *arXiv preprint arXiv:1705.04612* **2017**.
- (23) Kotsias, P.-C.; Arús-Pous, J.; Chen, H.; Engkvist, O.; Tyrchan, C.; Bjerrum, E. J. Direct steering of de novo molecular generation with descriptor conditional recurrent neural networks. *Nature Machine Intelligence* **2020**, *2* (5), 254-265.
- (24) Polykovskiy, D.; Zhebrak, A.; Sanchez-Lengeling, B.; Golovanov, S.; Tatanov, O.; Belyaev, S.; Kurbanov, R.; Artamonov, A.; Aladinskiy, V.; Veselov, M. Molecular sets (MOSES): a benchmarking platform for molecular generation models. *Frontiers in pharmacology* **2020**, *11*, 565644.
- (25) Prykhodko, O.; Johansson, S. V.; Kotsias, P.-C.; Arús-Pous, J.; Bjerrum, E. J.; Engkvist, O.; Chen, H. A de novo molecular generation method using latent vector based generative adversarial network. *Journal of Cheminformatics* **2019**, *11* (1), 1-13.
- (26) Arús-Pous, J.; Johansson, S. V.; Prykhodko, O.; Bjerrum, E. J.; Tyrchan, C.; Raymond, J.-L.; Chen, H.; Engkvist, O. Randomized SMILES strings improve the quality of molecular generative models. *Journal of cheminformatics* **2019**, *11* (1), 1-13.
- (27) Blaschke, T.; Arús-Pous, J.; Chen, H.; Margreitter, C.; Tyrchan, C.; Engkvist, O.; Papadopoulos, K.; Patronov, A. REINVENT 2.0: an AI tool for de novo drug design. *Journal of chemical information and modeling* **2020**, *60* (12), 5918-5922.
- (28) Kang, S.; Cho, K. Conditional molecular design with deep generative models. *Journal of chemical information and modeling* **2018**, *59* (1), 43-52.
- (29) Wang, J.; Hsieh, C.-Y.; Wang, M.; Wang, X.; Wu, Z.; Jiang, D.; Liao, B.; Zhang, X.; Yang, B.; He, Q. Multi-constraint molecular generation based on conditional transformer, knowledge distillation and reinforcement learning. *Nature Machine Intelligence* **2021**, *3* (10), 914-922.
- (30) Olivecrona, M.; Blaschke, T.; Engkvist, O.; Chen, H. Molecular de-novo design through deep reinforcement learning. *Journal of cheminformatics* **2017**, *9* (1), 1-14.
- (31) Krishnan, S. R.; Bung, N.; Bulusu, G.; Roy, A. Accelerating de novo drug design against novel proteins using deep learning. *Journal of Chemical Information and Modeling* **2021**, *61* (2), 621-630.
- (32) Guimaraes, G. L.; Sanchez-Lengeling, B.; Outeiral, C.; Farias, P. L. C.; Aspuru-Guzik, A. Objective-reinforced generative adversarial networks (organ) for sequence generation models. *arXiv preprint arXiv:1705.10843* **2017**.
- (33) Shi, C.; Xu, M.; Zhu, Z.; Zhang, W.; Zhang, M.; Tang, J. Graphaf: a flow-based autoregressive model for molecular graph generation. *arXiv preprint arXiv:2001.09382* **2020**.
- (34) Schneuing, A.; Du, Y.; Harris, C.; Jamasb, A.; Igashov, I.; Du, W.; Blundell, T.; Lió, P.; Gomes, C.; Welling, M. Structure-based drug design with equivariant diffusion models. *arXiv preprint arXiv:2210.13695* **2022**.

- (35) Igashov, I.; Stärk, H.; Vignac, C.; Satorras, V. G.; Frossard, P.; Welling, M.; Bronstein, M.; Correia, B. Equivariant 3d-conditional diffusion models for molecular linker design. *arXiv preprint arXiv:2210.05274* **2022**.
- (36) Gaines, B. B.; Bi, J. A deep molecular generative model based on multi-resolution graph variational Autoencoders. **2021**.
- (37) Button, A.; Merk, D.; Hiss, J. A.; Schneider, G. Automated de novo molecular design by hybrid machine intelligence and rule-driven chemical synthesis. *Nature machine intelligence* **2019**, *1* (7), 307-315.
- (38) Shen, C.; Krenn, M.; Eppel, S.; Aspuru-Guzik, A. Deep molecular dreaming: Inverse machine learning for de-novo molecular design and interpretability with surjective representations. *Machine Learning: Science and Technology* **2021**, *2* (3), 03LT02.
- (39) Gaudin, T.; Nigam, A.; Aspuru-Guzik, A. Exploring the chemical space without bias: data-free molecule generation with DQN and SELFIES. In *Second Workshop on Machine Learning and the Physical Sciences NeurIPS*, 2019.
- (40) Flam-Shepherd, D.; Wu, T. C.; Aspuru-Guzik, A. MPGVAE: improved generation of small organic molecules using message passing neural nets. *Machine Learning: Science and Technology* **2021**, *2* (4), 045010.
- (41) Nigam, A.; Pollice, R.; Aspuru-Guzik, A. Parallel tempered genetic algorithm guided by deep neural networks for inverse molecular design. *Digital Discovery* **2022**, *1* (4), 390-404.
- (42) Sanchez-Lengeling, B.; Aspuru-Guzik, A. Inverse molecular design using machine learning: Generative models for matter engineering. *Science* **2018**, *361* (6400), 360-365.
- (43) 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* (2), 268-276.
- (44) Griffiths, R.-R.; Hernández-Lobato, J. M. Constrained Bayesian optimization for automatic chemical design using variational autoencoders. *Chemical science* **2020**, *11* (2), 577-586.
- (45) Iwata, H.; Nakai, T.; Koyama, T.; Matsumoto, S.; Kojima, R.; Okuno, Y. VGAE-MCTS: a New Molecular Generative Model combining Variational Graph Auto-Encoder and Monte Carlo Tree Search. **2023**.
- (46) Ma, R.; Luo, T. PI1M: a benchmark database for polymer informatics. *Journal of Chemical Information and Modeling* **2020**, *60* (10), 4684-4690.
- (47) Ma, R.; Zhang, H.; Luo, T. Exploring high thermal conductivity amorphous polymers using reinforcement learning. *ACS Applied Materials & Interfaces* **2022**, *14* (13), 15587-15598.
- (48) 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. *Chemistry of Materials* **2021**, *33* (17), 7008-7016.
- (49) 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. *Chemistry of Materials* **2020**, *32* (24), 10489-10500.
- (50) 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* (24), 7669-7675.

- (51) Wildman, S. A.; Crippen, G. M. Prediction of physicochemical parameters by atomic contributions. *Journal of chemical information and computer sciences* **1999**, *39* (5), 868-873.
- (52) Kim, S.; Schroeder, C. M.; Jackson, N. E. Open Macromolecular Genome: Generative Design of Synthetically Accessible Polymers. *ACS Polymers Au* **2023**.
- (53) Zhang, J.; Mercado, R.; Engkvist, O.; Chen, H. Comparative study of deep generative models on chemical space coverage. *Journal of Chemical Information and Modeling* **2021**, *61* (6), 2572-2581.
- (54) Weng, G.; Zhao, H.; Nie, D.; Zhang, H.; Liu, L.; Hou, T.; Kang, Y. Rediscmol: Benchmarking molecular generation models in biological properties. *Journal of Medicinal Chemistry* **2024**.
- (55) Nigam, A.; Pollice, R.; Tom, G.; Jorner, K.; Willes, J.; Thiede, L. A.; Kundaje, A.; Aspuru-Guzik, A. Tartarus: A benchmarking platform for realistic and practical inverse molecular design. *arXiv preprint arXiv:2209.12487* **2022**.
- (56) Kim, S.; Chen, J.; Cheng, T.; Gindulyte, A.; He, J.; He, S.; Li, Q.; Shoemaker, B. A.; Thiessen, P. A.; Yu, B. PubChem 2019 update: improved access to chemical data. *Nucleic acids research* **2019**, *47* (D1), D1102-D1109.
- (57) Wang, Y.; Xiao, J.; Suzek, T. O.; Zhang, J.; Wang, J.; Bryant, S. H. PubChem: a public information system for analyzing bioactivities of small molecules. *Nucleic acids research* **2009**, *37* (suppl_2), W623-W633.
- (58) Eyben, F.; Wöllmer, M.; Schuller, B. Opensmile: the munich versatile and fast open-source audio feature extractor. In *Proceedings of the 18th ACM international conference on Multimedia*, 2010; pp 1459-1462.
- (59) Radford, A.; Narasimhan, K.; Salimans, T.; Sutskever, I. Improving language understanding with unsupervised learning. **2018**.
- (60) Mercado, R.; Rastemo, T.; Lindelöf, E.; Klambauer, G.; Engkvist, O.; Chen, H.; Bjerrum, E. J. Graph networks for molecular design. *Machine Learning: Science and Technology* **2021**, *2* (2), 025023.
- (61) Mercado, R.; Rastemo, T.; Lindelöf, E.; Klambauer, G.; Engkvist, O.; Chen, H.; Bjerrum, E. J. Practical notes on building molecular graph generative models. *Applied AI Letters* **2020**, *1* (2).
- (62) Benhenda, M. ChemGAN challenge for drug discovery: can AI reproduce natural chemical diversity? *arXiv preprint arXiv:1708.08227* **2017**.
- (63) Mayr, A.; Klambauer, G.; Unterthiner, T.; Steijaert, M.; Wegner, J. K.; Ceulemans, H.; Clevert, D.-A.; Hochreiter, S. Large-scale comparison of machine learning methods for drug target prediction on ChEMBL. *Chemical science* **2018**, *9* (24), 5441-5451.
- (64) Bento, A. P.; Gaulton, A.; Hersey, A.; Bellis, L. J.; Chambers, J.; Davies, M.; Krüger, F. A.; Light, Y.; Mak, L.; McGlinchey, S. The ChEMBL bioactivity database: an update. *Nucleic acids research* **2014**, *42* (D1), D1083-D1090.
- (65) Irwin, J. J.; Sterling, T.; Mysinger, M. M.; Bolstad, E. S.; Coleman, R. G. ZINC: a free tool to discover chemistry for biology. *Journal of chemical information and modeling* **2012**, *52* (7), 1757-1768.
- (66) Putin, E.; Asadulaev, A.; Vanhaelen, Q.; Ivanenkov, Y.; Aladinskaya, A. V.; Aliper, A.; Zhavoronkov, A. Adversarial threshold neural computer for molecular de novo design. *Molecular pharmaceutics* **2018**, *15* (10), 4386-4397.
- (67) Sterling, T.; Irwin, J. J. ZINC 15—ligand discovery for everyone. *Journal of chemical information and modeling* **2015**, *55* (11), 2324-2337.

