

# Polymers for Extreme Conditions Designed Using Syntax-Directed Variational Autoencoders

Rohit Batra,\* Hanjun Dai, Tran Doan Huan, Lihua Chen, Chiho Kim, Will R. Gutekunst, Le Song, and Rampi Ramprasad



Cite This: <https://dx.doi.org/10.1021/acs.chemmater.0c03332>



Read Online

ACCESS |



Metrics & More




Article Recommendations



Supporting Information

**ABSTRACT:** The design/discovery of new materials is highly nontrivial owing to the near-infinite possibilities of material candidates and multiple required property/performance objectives. Thus, machine learning tools are now commonly employed to virtually screen material candidates with desired properties by learning a theoretical mapping from material-to-property space, referred to as the forward problem. However, this approach is inefficient and severely constrained by the candidates that the human imagination can conceive. Thus, in this work on polymers, we tackle the materials discovery challenge by solving the inverse problem: directly generating candidates that satisfy desired property/performance objectives. We utilize syntax-directed variational autoencoders (VAE) in tandem with Gaussian process regression (GPR) models to discover polymers expected to be robust under three extreme conditions: (1) high temperatures, (2) high electric field, and (3) high temperature and high electric field, useful for critical structural, electrical, and energy storage applications. This approach to learn from (and augment) human ingenuity is general and can be extended to discover polymers with other targeted properties and performance measures.



## INTRODUCTION

A range of materials science problems are now being tackled using machine learning (ML), including discovery of materials,<sup>1–3</sup> revelation of hidden structure–property relations,<sup>4–6</sup> recommendations on plausible materials or chemical synthesis pathways,<sup>7</sup> and design of automated experiments.<sup>8</sup> Arguably, the most direct way ML has influenced materials science is by allowing virtual screening of thousands or millions of new materials with desired functionalities,<sup>9,10</sup> which is far beyond the current capabilities of commonly employed high-throughput simulation or empirical techniques. This paradigm shift has resulted in the discovery of superhard ceramics,<sup>11</sup> metallic glasses,<sup>12</sup> superconducting alloys,<sup>13</sup> photovoltaics,<sup>14,15</sup> redox flow batteries,<sup>16</sup> among others, and continues to revolutionize various aspects of materials and chemical sciences.<sup>17–20</sup>

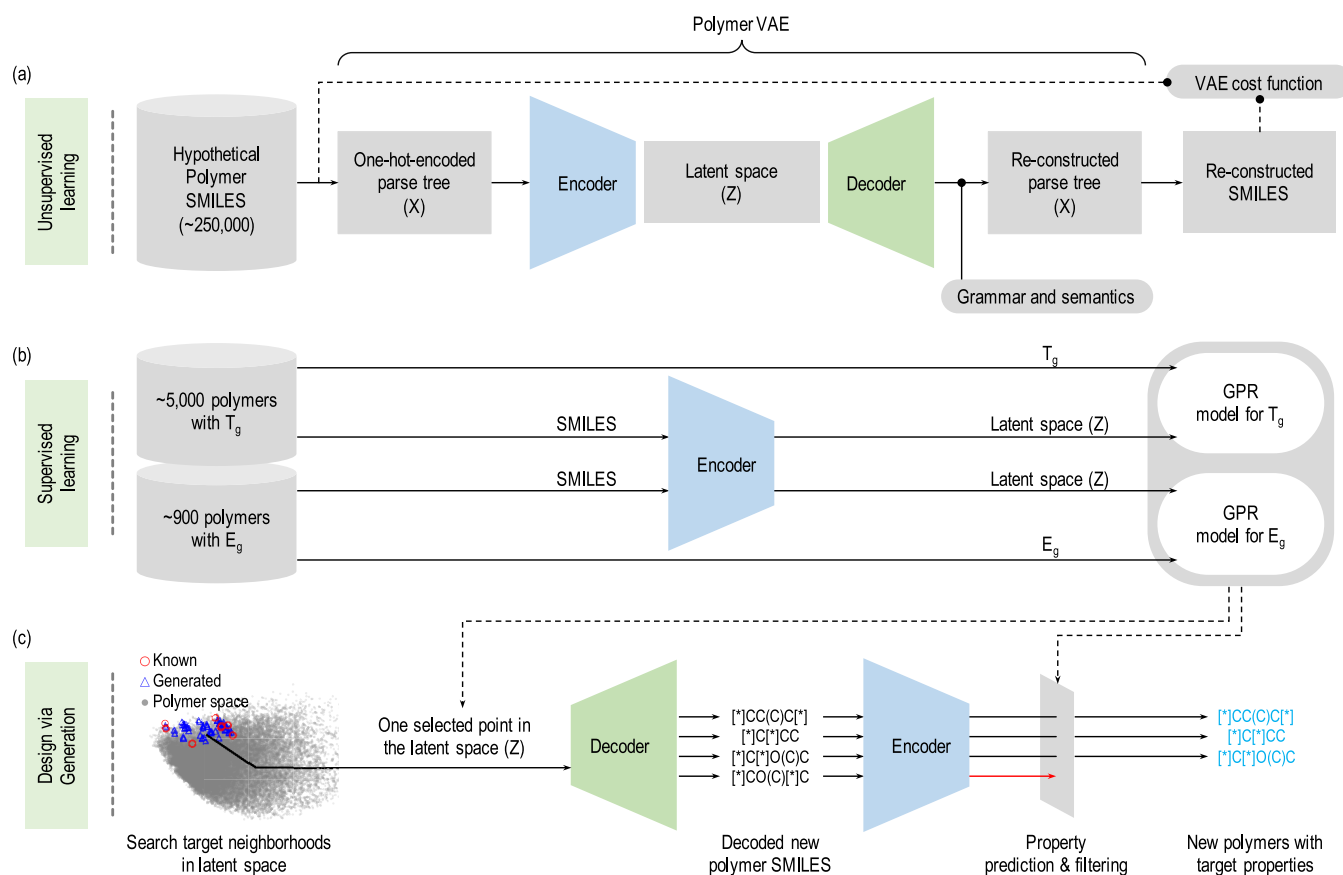
The materials discovery process has, however, been mostly dominated by ML techniques that attempt to solve the forward problem: “given a material, what is its property?” Under this, starting from a dataset of materials and their corresponding properties, supervised learning models are trained, using which property predictions for a large pool of plausible materials are made and candidates with desirable properties are down-selected—to be later verified by experiments. Cost-effective techniques, such as Bayesian optimization<sup>21</sup> and active learning,<sup>22</sup> are often utilized to search over relatively larger chemical space. Nonetheless, these property-prediction

approaches are restricted to (and implicitly biased toward) the initial pool of materials considered, typically falling well within the purview of human imagination. This is a major limitation considering the near-infinite space offered by materials—a realistic estimate being placed around  $10^{100}$ , which is further compounded if one considers complications that arise due to the hierarchy of scale, stretching from sub-nanometer to microscopic and mesoscopic.<sup>23</sup>

Ideally, the problem of materials discovery is of inverse type: starting with a set of desired properties, the aim is to search for materials that display those functionalities. In this regard, two distinct schemes have recently been advanced. On the one hand, global search strategies such as genetic algorithm,<sup>24</sup> evolutionary searches, etc.<sup>15</sup> tackle this inverse problem by iteratively mutating/mixing material candidates until the target properties are met (or a cost function is minimized); on the other hand, deep generative ML models provide a more straightforward solution by learning an inverse functional mapping from a latent to the materials space. Although both these approaches allow us to go beyond the constraints of

Received: August 14, 2020

Revised: November 21, 2020



**Figure 1.** Overall strategy to design polymers with targeted properties by solving the inverse problem of properties-to-materials space. (a) First, a variational autoencoder is used to map polymer SMILES to and from a continuous latent space. (b) Polymers with known properties ( $T_g$  and  $E_g$  here) are then encoded (or “fingerprinted” in the latent space) using the encoder. A supervised learning technique (GPR here) is then used to map these fingerprints to the different polymer properties, resulting in ML property models. (c) Finally, in the design stage, known polymers with desirable properties are first encoded to find the region of interest in the latent space. Using ML property-prediction models, those latent points are sampled that satisfy a given design goal. The decoder is utilized next to construct the corresponding polymer SMILES, which goes through a subsequent round of encoding and ML prediction to ensure that the decoded polymers meet the design goals.

human imagination, the latter additionally learns inherent chemical trends present in the materials data and allows materials design by interpolating between known suitable materials or searching within the vicinity of a known best material.

Generative models, such as variational autoencoder (VAE)<sup>25–29</sup> and generative adversarial network (GAN),<sup>30–32</sup> have been recently used to design drug-like molecules<sup>33,34</sup> or even polymers<sup>35</sup> and periodic solids.<sup>36</sup> For instance, Gómez-Bombarelli et al.<sup>25</sup> represented the chemical structure of molecules using the SMILES (or simplified molecular-input line-entry system) language, where the encoder part of the VAE learns to project the molecular SMILES onto the latent space, while the decoder unit reconstructs the molecular SMILES given its encoded latent representation. However, a key challenge in developing such a generative model for molecules (or for materials, in general) is to learn the mapping between a continuous latent space and a discrete chemical structure space, enabled by the syntax or grammar of the SMILES language and the semantics associated with the chemistry of molecules (or materials). The absence of explicit syntax and semantics in the VAE model allows unnecessary flexibility in representing SMILES and thus requires more training data to learn the different chemical rules underlying chemical structures (and the SMILES language). This can also

result in the poor quality of the learned latent space and a high occurrence of invalid SMILES upon decoding, making the process of molecular discovery inefficient. A fundamental improvement was achieved when Kusner et al.<sup>26</sup> successfully incorporated the structure or syntax of SMILES language using the ideas of context-free-grammar. They used grammar associated with the SMILES language to convert molecular SMILES to parse trees, one-hot representation of which were then used to train the VAE models. The final missing piece of incorporating semantics (e.g., respecting elemental valency) was recently achieved by Dai et al.<sup>27</sup> through the concept of stochastic lazy attributes, with the overall model termed as syntax-directed VAE. Inclusion of both syntax and semantics was shown to drastically improve the performance of the VAE by learning a more meaningful latent representation, but more importantly, it allowed decoding of—syntactically and semantically—valid molecular SMILES at a significantly higher rate, making feasible the use of generative models for materials discovery. Recently, a very powerful and general approach, titled SELFIES, has been shown to outperform above-mentioned methods by incorporating both syntactic and semantic constraints through derivation rules table.<sup>34</sup> However, extension of this method for the case of polymers, requiring special constraints, such as the presence of chain ends with matching bond type, remains to be established. The

concepts of syntax and semantics are not limited to materials science and find applications in other areas, such as symbolic expressions<sup>37</sup> and programming languages.<sup>38</sup>

In this work, we go beyond molecular design and develop polymer-specific syntax-directed VAE to discover polymers with extreme thermal and electrical stability. Although VAE has been used for polymer discovery in previous works, they are either limited by the explored chemical space due to the involved grammar being restricted to a predetermined set of chemical fragments<sup>35</sup> or rely on generation of large molecules as proxy to polymers.<sup>39,40</sup> In contrast, this work on polymer generation uses VAE that incorporates both syntax as well as semantics constraints. To achieve this, crucial modifications were introduced in the SMILES grammar, in addition to inclusion of polymer-specific semantics (e.g., 2 chain ends). Thus, the decoding process incorporates polymer syntax and semantics using context-free-grammar parse trees and stochastic lazy attributes to ensure that decoded SMILES represents plausible polymers. Another critical challenge to train a polymer VAE, as opposed to other studies on molecular VAE, is the availability of large amounts of training examples (in the order of hundred thousands). The polymer community is particularly known to suffer from data sparsity problems. For instance, to-date, the total number of chemically diverse polymers synthesized is roughly  $\sim 12\,000$ , which is not enough to learn a good-performing VAE model; in contrast, datasets consisting of over 250 000 molecules are available. Here, we resolve this issue using retro-synthetic ideas to generate a representative hypothetical dataset of  $\sim 250\,000$  polymers, constructed from a few thousand molecular 'building blocks' or fragments derived from  $\sim 12\,000$  empirically synthesized polymers obtained from the literature.

To demonstrate the preeminence of the current approach for polymer design, we use it to discover several hundred polymer candidates that are expected to perform well, and be stable, under 3 extreme conditions: (1) high temperatures, (2) high electric field, and (3) high temperature and high electric field. We make the further assumptions that high temperature and high electric field behaviors may be related to the polymer glass transition temperature ( $T_g$ ) and bandgap ( $E_g$ ), respectively. In particular, we set our property goals to be (1)  $T_g > 600$  K, (2)  $E_g > 6.5$  eV, and (3)  $T_g > 500$  K and  $E_g > 4$  eV, pertaining to the three types of extreme conditions we consider here. While high  $T_g$  polymers are desired for their mechanical robustness at elevated temperatures,<sup>41–44</sup> polymers with high  $E_g$  offer a large range of electrical stability at low dielectric loss and are useful for high-energy capacitor dielectrics.<sup>45,46</sup> The last design goal, namely, high temperature and high electric field stability, is especially important for high energy density dielectrics that are both thermally stable and display high electrical breakdown strength. The occurrence of inverse relation between the  $T_g$  and  $E_g$  of polymers makes the last goal most difficult to achieve from a materials discovery perspective.

Figure 1 captures the overall strategy used here for polymers design. First, an unsupervised learning task is performed to train the VAE. Critical aspects of this process include the conversion of polymer SMILES to and from one-hot-encoded parse trees and incorporation of grammar and semantics directly into the decoding step. The latter regulates the probability space such that only cases with valid SMILES have nonzero probability of generation. It should be noted that the probabilistic formulation of the latent variable in the

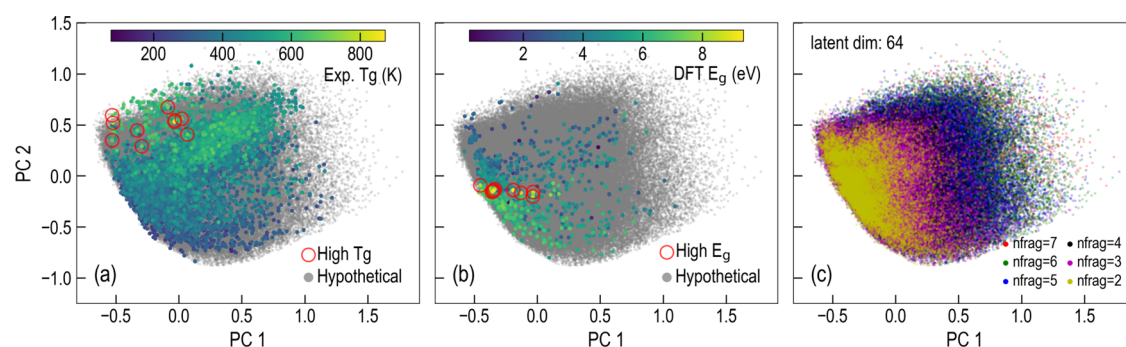
autoencoder introduces stochasticity in the encoding and the decoding process, resulting in a distribution of polymer SMILES associated with a particular latent vector and vice-versa. Next, in the supervised learning stage, polymers with known properties— $E_g$  and  $T_g$  here—are first encoded in their latent representation, using which Gaussian process regression (GPR) models are trained for property prediction (cf. Figure 1). Chemically meaningful latent space learned by the VAE helps build accurate GPR models for property predictions. Finally, to design new polymers with a given set of target properties, a simple enumeration followed by a generative interpolation approach is applied. Under this, known (or hypothetical) polymers that satisfy the given design criteria are initially encoded to find regions in the latent space where desirable polymers are expected to be present. Linear interpolations within these desired regions of the latent space are then used to select latent points for which GPR property predictions meet the desired goals. The polymer SMILES associated with such selected latent vectors is decoded, followed by another round of encoding and GPR property predictions to ascertain that the designed polymer indeed satisfies the design objectives. The full design process is illustrated under the heading of "Design via Generation" in Figure 1.

Finally, we note that the presented scheme is general and could be used to design polymers with any set of property targets (as long as predictors for the desired set of properties are available). By establishing an inverse mapping from latent representation to polymer SMILES, it overcomes the limitations of the property-prediction (or forward problem based) approaches that screen polymers from a predetermined dataset and suffers from selection bias. Moving forward, we expect this scheme to be further refined using the concepts of transfer learning, multitask learning, and semisupervised learning.

## RESULTS AND DISCUSSION

**Polymer Variational Autoencoder.** As stated earlier, the VAE is a type of generative model that learns to encode or decode a sequence of inputs to and from a continuous latent space. This is achieved using deep neural networks (feed-forward, convolution or recurrent type) to represent the encoding and the decoding units, with the constraint that the encoder lowers the dimensionality of (or compresses) the input data, while the decoder performs the decompression operation to reconstruct the input sequence (see Figure 1a). As an outcome of this data compression and expansion exercise, the VAE is expected to learn essential features of the data—polymer chemistry in our case. Here, we train the VAE to map the sequence of polymer SMILES to and from a continuous latent representation, allowing us to search this latent space for polymers with desired properties. Notably, VAE is an unsupervised learning method that requires just the polymer SMILES for training, without any associated property information (labeled data). Given the dearth of property data in materials science, this is a significant advantage.

Crucially, the VAE is constructed such that it incorporates both the syntax of the SMILES language and the semantics of the polymer chemistry. This is achieved using context-free-grammar parse trees and stochastic lazy attributes, details of which can be found in the section **Methods**. In particular, we modified the SMILES grammar and incorporated new attributes to ensure that the generated polymers have exactly



**Figure 2.** Learned latent space of the hypothetical SMILES dataset as projected onto the first two principal components. Observed trends in the (a) empirical  $T_g$  and the (b) density functional theory (DFT)-computed  $E_g$  values when the known polymer datasets are encoded into the same latent space, with the highest 10 property value polymers clustered in a limited region (red color symbols). (c) Increase in the polymer diversity with increasing “nfrag” or number of molecular building blocks (see Methods) as reflected by the area spanned in the latent space.

two chain ends with the same bonding type (single, double bond, etc.) Other syntax and semantics constraints that are also applicable to molecules, such as valency of different elements, closure of rings, etc., were also retained (see Methods). Henceforth, we will refer to this syntax-directed VAE as just VAE. The inclusion of grammar and semantics frees the VAE from explicitly (or spending extra efforts in) learning such rules from the training set and allows it to wholly concentrate on finding more important chemical trends, making the latent representation more chemically meaningful. Further, it ensures the generation (or decoding) of valid polymers at a significantly higher rate, making the process of polymer design computationally efficient and practically feasible. In this work, two VAE models with latent dimensions 64 and 128, respectively, termed as VAE-64 and VAE-128, were trained using a hypothetical polymer SMILES dataset (with  $\sim 250\,000$  SMILES). Details on the VAE training, including the hypothetical SMILES dataset and the model architecture, are discussed in the section Methods.

Figure 2 captures the ability of the polymer VAE to learn a chemically meaningful latent representation through just the unsupervised learning of SMILES. In panels (a) and (b), gray colored symbols mark the latent projections of the entire hypothetical SMILES dataset, reduced into the first two principal components. Likewise, the polymers with known  $T_g$  and  $E_g$ , respectively, named as  $T_g$  and  $E_g$  datasets, are converted to their latent representations, followed by their projection onto the same two principal components. More details on these two property datasets, including their source and data size, can be found in the section Methods. Clear trends on the concentration of polymers with similar  $T_g$  and  $E_g$  in Figure 2a,b suggest that the VAE has indeed learned important chemical characteristics. Moreover, red colored symbols denoting the top 10 known polymers with the highest  $T_g$  and  $E_g$  values appear in a small region of the latent space, highlighting the area where similar polymers with perhaps even higher  $T_g$  and  $E_g$  values can be found. The presence of these chemical trends can be explained based on the following rationale: during training, VAE learns to place polymers with similar SMILES in the neighboring regions of the latent space, thereby automatically learning the chemical information present in the SMILES representation. We contrast this observation with the previous study on molecules, where Gómez-Bombarelli et al.<sup>25</sup> were unable to see such clear chemical trends. Possibly, this is because their VAE model does not respect SMILES grammar and molecular semantics, as was

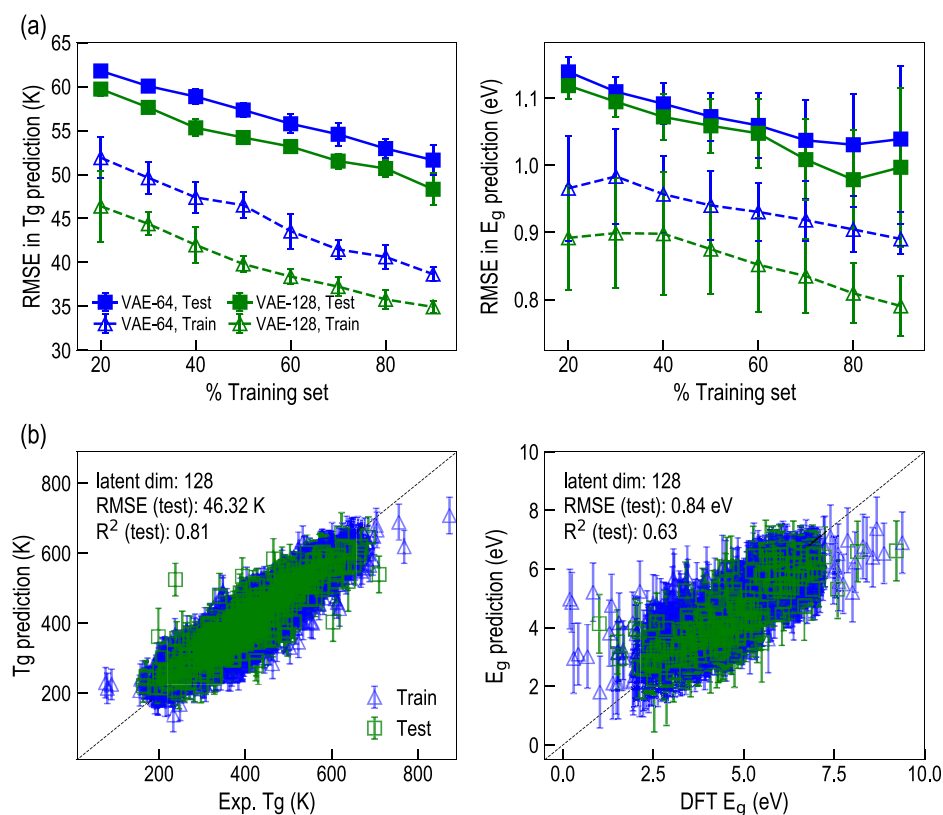
discussed in the recent works.<sup>26</sup> Finally, in panel (c) we demonstrate how the diversity of the hypothetical SMILES dataset increases with the increase in the number of molecular building blocks chosen to create the polymers. While Figure 2 presents results for VAE-64, similar observations using VAE-128 are provided in Supporting Information (SI), Figure S1.

Other metrics to measure the performance of a VAE are the reconstruction accuracy and the prior validity. The former captures the ability of the model to reconstruct polymer SMILES from their corresponding latent representations, a measure of the quality of the latent space; the latter measures the probability of the decoder to generate valid SMILES when sampling from the prior latent space and quantifies the computational efficiency of the decoding process. Considering the probabilistic nature of the encoder and the decoder, the reconstruction accuracy was computed by encoding each polymer 10 times, followed by decoding 100 times. This results in 1000 decoded SMILES for each input polymer. Two types of reconstruction accuracy were defined: (1) strict, when the most frequent of these 1000 decodings were the same as the input polymer, and (2) loose, when at least one of the 1000 decodings were the same as the input polymer. The prior validity was computed by sampling 1000 latent points from the prior distribution  $p(\mathbf{z}) = \mathcal{N}(0, \mathbf{I})$ , decoding each of these points 500 times, and finding the overall percentage of valid SMILES generated. Good performance for the average reconstruction accuracy for the  $T_g$  and the  $E_g$  datasets and the percentage prior validity for the case of VAE-64 and VAE-128 can be seen in Table 1. We note that both  $T_g$  and  $E_g$  datasets were not included during the VAE training and

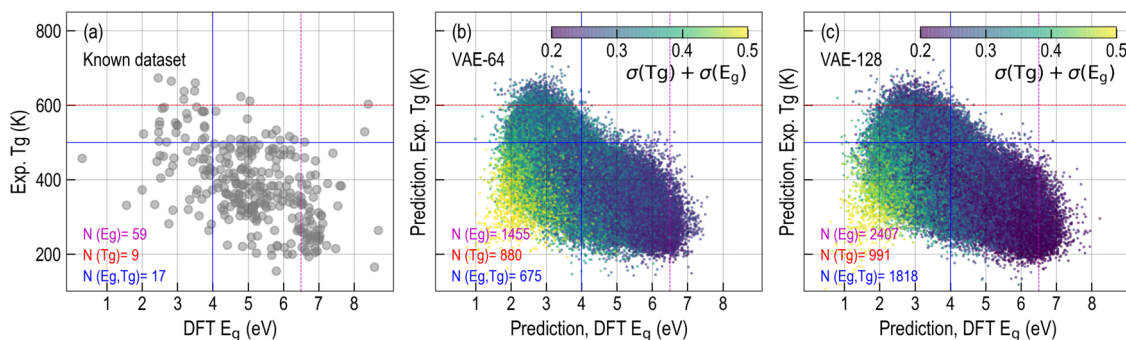
**Table 1. Performance of the Polymer VAE as a Function of Latent Dimensions for the  $T_g$  and  $E_g$  Datasets<sup>a</sup>**

model	prior validity (%)	reconstruction (%)		ML $R^2$ (RMSE)	
		$T_g$ (K)	$E_g$ (eV)	$T_g$ (K)	$E_g$ (eV)
VAE-64	27	51 (75)	53 (86)	0.76 (51.63)	0.50 (1.04)
VAE-128	13	57 (78)	67 (89)	0.79 (48.32)	0.53 (1.03)

<sup>a</sup>See the text for the details on the definitions of two types of reconstruction accuracy, strict and loose (in brackets), and prior validity.  $R^2$  coefficient and RMSE on the test set (20 % of dataset) for the two ML (GPR) models are also provided.



**Figure 3.** (a) Learning curves for the GPR models built using the  $T_g$  and  $E_g$  datasets. The polymers were first encoded using VAE, which was then mapped to their corresponding  $T_g$  or  $E_g$  property values. (b) Example parity plots demonstrating the performance of the GPR  $T_g$  and  $E_g$  models built using VAE-128 latent representation.



**Figure 4.** Comparison of the trend between  $T_g$  and  $E_g$  for (a) known polymers and for hypothetical polymers as predicted using GPR models based on latent representations of (b) VAE-64 and (c) VAE-128. In (b) and (c), the marker color represents the sum of the uncertainties in GPR  $T_g$  and  $E_g$  predictions, each normalized by its range. The magenta, red, and blue colored lines are the guide to the three design criteria high  $E_g$ , high  $T_g$ , and high  $E_g$  and  $T_g$ , respectively, with  $N(E_g)$ ,  $N(T_g)$ , and  $N(E_g, T_g)$  highlighting the number of polymers found to satisfy the three design criteria (in the same order) in different datasets.

constitute unseen “new” polymer cases. While the reconstruction accuracy of VAE-128 is better than VAE-64, perhaps owing to larger latent dimensionality or capacity to retain information, its performance is slightly worse in terms of prior validity. A fair comparison between the VAE results presented here is not possible with the past works on molecules owing to the different dataset types and variations in chemical diversity. However, we note that a similar level of performance was seen in the past for the case of molecules; reconstruction accuracy and prior validity were, respectively, reported to be 44.6 and 0.7 without incorporating grammar and semantics,<sup>25</sup> 53.7 and 7.2 with grammar only,<sup>26</sup> and 76.2 and 43.5 with both grammar and semantics.<sup>27</sup>

Having established the validity of the latent space learned by polymer VAE, we next build GPR models using the  $T_g$  and  $E_g$  datasets. For this, the polymer latent representations obtained from the VAE model were mapped to their respective property values (see Methods). The learning curves for the two properties using both the VAE-64 and VAE-128 latent representations are provided in Figure 3a. In all cases, reasonable error trends with the test error decreasing with increasing training set size can be seen. Provided  $T_g$  dataset is estimated to have a noise of around 30 K, reaching test errors around 50 K is quite respectable. The performance of  $E_g$  models, although acceptable, is comparatively worse. This could be because of the larger inherent noise present in the

theoretically computed  $E_g$  values of polymers, especially with the assumption of treating them as single chains (see details on the source of  $E_g$  dataset in the [Methods](#) section). In corroboration with the results obtained for the reconstruction accuracy, the performance of VAE-128 GPR models is visibly better than VAE-64 for both the property datasets. Example parity plots illustrating the performance of the GPR property-prediction models, using VAE-128 latent representations, are shown in [Figure 3b](#). Two quantitative error measures, i.e., root mean square error and correlation coefficient, averaged over 10 statistical runs, are also provided in [Table 1](#) for both the datasets. Once again, we note that the good level of accuracy reached by the GPR models is a direct indicator of the high quality of latent representation learned by the VAE.

**Polymer Design.** As example problems, we choose to design three classes of polymers with the following property objectives: (1)  $T_g > 600$  K, (2)  $E_g > 6.5$  eV, and (3)  $T_g > 500$  K and  $E_g > 4$  eV. Polymers with these functionalities are desirable for various dielectric applications; high  $T_g$  polymers provide high mechanical strength over larger temperature window,<sup>44</sup> high  $E_g$  polymers intrinsically have low dielectric loss and high enough electrical breakdown strength,<sup>47</sup> and high  $T_g$  and  $E_g$  polymers are expected to be great candidates for high energy density capacitor applications, particularly, at higher temperatures. While at low temperatures, bandgap has been found to correlate well with the breakdown strength of polymers,<sup>47</sup> at high temperatures other factors, such as chain-packing or defects, may become more relevant. Thus, high  $E_g$  as a criterion for the design of high breakdown strength polymers is only approximate and can be replaced with more accessible or correlated properties. The difficulty level of these three design goals becomes apparent from the  $T_g$  vs  $E_g$  plot in [Figure 4a](#), wherein all of the 303 known polymers common to both the property datasets are drawn. The number of polymers known to satisfy the aforementioned 3 design goals, i.e., high  $T_g$ , high  $E_g$ , and high  $T_g$  and  $E_g$ , are 9, 59, and 17, respectively. Further, an inverse relation between the  $T_g$  and  $E_g$  illustrates the difficulty of the last design goal. Finally, we note that our aim here is to find chemically diverse polymers that satisfy these design goals rather than finding polymers that maximize/minimize a particular property. This is relevant from an application standpoint because for a successful dielectric, a polymer must satisfy a collection of properties, including high charge transport barrier, high breakdown strength, low dielectric loss, etc., rather than display one best property. By imposing many such simple design criteria, the hope is that several polymer candidates can be proposed in extensions to this work, which can be later down-selected by imposing other screening criteria based on a target application.

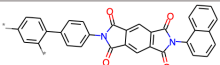
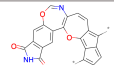
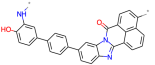
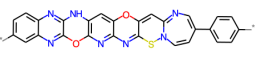
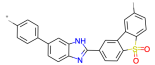
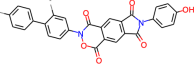






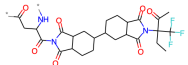
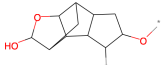
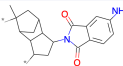
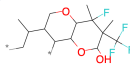
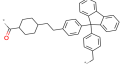
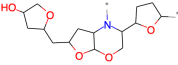
**Design via Enumeration.** As part of the first design problem, we simply make property predictions for the entire hypothetical SMILES dataset of nearly 250 000 polymers. The  $T_g$  and  $E_g$  GPR models were used to make these predictions based on the VAE-encoded latent representations. The inverse relation between the  $T_g$  and  $E_g$  as observed for the case of known polymers in [Figure 4a](#) is also preserved when using the GPR property predictions for the hypothetical polymer dataset in [Figure 4b,c](#). This, along with test set accuracy reached by the GPR models, provides some confidence in the property predictions made for the hypothetical polymer dataset using the present scheme. There are, however, some polymers with high uncertainty in the GPR predictions, as captured by the marker color codings in panels (b) and (c). Nonetheless, the

yellow colored high uncertainty regions are different from those where the three design criteria are met. The number of hypothetical polymers that satisfy the three design criteria are also included in panels (b) and (c), which is significantly larger than that in the original  $T_g$  and  $E_g$  datasets. The complete list of these new polymer candidates predicted to have the desired functionalities and designed through solving the forward problem is provided in the [SI](#).

**Design via Inversion.** The chemical trends captured by the latent representation, illustrated in [Figure 2](#), suggest that polymers with similar properties occupy a closer neighborhood in the latent space. Thus, if a few polymers with the desired properties are already known, a search in their latent neighborhood could reveal many more suitable polymer candidates. The GPR property prediction models could be used to advance this search. Once the latent points satisfying the given design criteria are identified, the inverse mapping from the latent representation to the polymer, made available through the VAE decoder, can be used to find the associated polymer candidate. We note that owing to the stochastic nature of the decoding process, there is some ambiguity with the exact polymer represented by the selected latent point. Thus, to address this issue, a second round of encoding for all of the decoded polymer candidates is performed and the associated GPR property predictions are made to check if the design criteria are indeed met. Polymers satisfying the design criteria are retained.

Following the above-mentioned design strategy, for a particular design goal, the top 250 polymer candidates identified as part of the earlier enumeration scheme were selected. Linear interpolations in the latent space for all combinations of selected 250 polymers were performed, and latent points whose associated GPR predictions meet the design criteria were selected. Owing to computational cost, decoding for only top 10 000 latent points—with the highest GPR property prediction values—was done 100 times. Only those decoded candidates were retained that satisfy the design criteria after another round of encoding, as discussed above and illustrated in [Figure 1c](#). This procedure was repeated for all three design criteria. The complete list of polymer candidates generated using this approach is provided in the [SI](#), while some representative cases are presented in [Figure 5](#). Care was taken to eliminate duplicate candidates with the  $E_g$ ,  $T_g$ , and the hypothetical SMILES datasets.

A few observations can be made from the results presented in [Figure 5](#). First, from a chemical standpoint the proposed polymer candidates appear to be rational; predominance of conjugated systems is reflective of polymers with high  $T_g$ , the presence of saturated C–F bonds suggest high  $E_g$ , and a balance between saturated rings and C bonded with highly electronegative F and O species is expected to have both high  $T_g$  and  $E_g$ . Though synthesis of these candidates and their subsequent characterization and measurement would be required to confirm their property predictions. Nonetheless, for the case  $E_g$  predictions we provide preliminary validation using DFT computations. We selected 30 new polymers (10 from each of the design goals) proposed using VAE design scheme and constructed single-chain polymer systems on which DFT calculations were performed, consistent with generation of the  $E_g$  training dataset. The results are similar to those observed for the test set in [Figure 3](#) and are provided in [SI](#), [Figure S2](#).

Example New Polymers		
#	Enumeration	Generation
High $T_g$ (in K)		
1	 686.77(53.78)	 737.76 (57.30)
2	 661.90(54.18)	 708.59(70.13)
3	 649.04(52.20)	 707.83(53.64)
High $E_g$ (in eV)		
1	 8.23 (1.04)	 8.61 (1.10)
2	 7.96 (1.02)	 8.56 (1.06)
3	 7.60 (1.02)	 8.46 (1.01)
High $T_g$ (in K), $E_g$ (in eV)		
1	 517.44 (81.82), 4.97 (1.67)	 509.53 (54.01), 6.32 (1.11)
2	 501.54 (65.21), 5.49 (1.34)	 528.65(56.11), 6.05 (1.19)
3	 548.52 (56.93), 4.61 (1.30)	 506.89 (52.82), 6.2 (1.11)

**Figure 5.** Example polymers designed in this work, along with their GPR  $T_g$  and  $E_g$  estimates and uncertainties (in brackets). The symbol “\*” represents the polymer chain ends.

Second, following the simple strategy of search in the latent space neighborhood, the generative or inverse scheme is indeed able to identify several cases, which have superior properties (higher  $T_g$  and/or  $E_g$ ) than that obtained through enumeration or property-prediction (forward problem) approach (also see the SI). This clearly highlights the advantage of the design strategy presented here (which allows us to go beyond the constraints of human imagination) and the motivation to solve the inverse problem. In contrast to the 17 known polymers, almost 300 new potential candidates have been designed (list provided in the SI) with the goal of exhibiting high  $T_g$  and  $E_g$ . Further, this scheme has the potential to find many more polymers satisfying these criteria. Finally, associated with each distinct design goal, polymers with very different chemistry are generated by the decoder. An analysis of the common chemical fragments present in the VAE-generated polymers for the three target properties is provided in the SI. It not only shows the distinct chemistry

necessary to satisfy each of the three conditions but also reveals new trends in the form of chemical fragments promoting these properties. For instance, saturated 4-, 5-membered rings, bridged rings, oxolane groups, and C atoms bonded to a limited number of electronegative F, Cl, and O groups had a higher occurrence rate for polymers predicted to have high  $T_g$  and  $E_g$ . Furthermore, since the VAE-generated polymers contain new chemical fragments as opposed to those present in the training data, it demonstrates the ability of our approach to explore extremely huge chemical space, without being restricted to a predetermined set of chemical fragments. Overall, the VAE approach resulted in generation of reasonable polymers, with distinct chemistry for different conditions, involving generation of new chemical fragments and with superior property predictions. We expect the inclusion of grammar and semantics in the VAE as the major cause that enabled learning of a meaningful latent space, making it a powerful polymer-discovery tool.

Many avenues remain for improving the design strategy adopted here. The set of polymer semantics included is not complete. Incorporating more semantics is expected to enhance the performance of the entire design process. The model provides no information regarding the synthesizability of the generated polymers; this point is particularly relevant for high  $T_g$  polymers that are difficult to synthesize. Moreover, rather than individually conducting the unsupervised (autoencoder) and the supervised (GPR) training, joint training can be performed for all of the SMILES as well as the property datasets. Such semisupervised learning has been shown to improve the performance of machine learning models. Efforts to incorporate other aspects of polymer structure, such as cross-linking, chain length, dispersity, etc., should also be made, perhaps through the use of recently developed BigSMILES representation.<sup>48</sup> Further, polymer search in the latent space could be potentially improved using reinforcement learning.<sup>49</sup>

## CONCLUSIONS

The space of potential materials is near-infinite, making the search for materials with desirable properties intractable. With the introduction of machine learning, the design of new materials with desirable properties has accelerated and/or improved. However, still most of the methods rely on solving the forward problem, wherein materials-to-property mappings are learned using which a predetermined list of material candidates is screened. This approach is not only limited in terms of the explored materials space but is computationally inefficient. Here, instead we demonstrate a strategy that solves the inverse problem: starting from the properties required, desirable materials candidates are generated. As an example, we design 3 classes of polymers applicable for thermal extremes, electrical extremes, and thermal and electrical extremes.

Our strategy is based on the use of variational autoencoders to encode and decode polymers to and from a continuous latent space, allowing one to search this space for new polymers with desirable properties. The autoencoder is trained to operate using the SMILES representation of the polymers. Crucially, the learning ability of the autoencoder is enhanced by directly incorporating the SMILES grammar and polymer semantics using context-free-grammar parse trees and stochastic lazy attributes. This allows to construct a more meaningful latent representation for polymers and improves the efficiency of the decoder to generate valid polymers. Next,

supervised learning GPR models were built for two property datasets, i.e.,  $T_g$  and  $E_g$ , using the polymer latent representations.

To design polymers with desired  $T_g$  and  $E_g$  properties, enumeration and generative schemes were adopted. In the former case, GPR property predictions for a hypothetical (but realistic) polymer dataset were made to screen polymers that meet the design criteria, parallel to solving the forward problem. These selected polymer candidates helped drive the latter case by identifying interesting latent neighborhoods, which were searched using GPR property prediction models to find latent points that correspond to the required properties. The decoder was then used to find the associated polymers. More importantly, among the two approaches, the generative design approach (which solves the inverse problem) was found to discover polymer candidates with superior properties, with a much lower computational budget. Using both these schemes, we proposed several hundred potential polymers for different dielectric applications, although we note that the proposed design scheme is quite general and can be applied to discover polymers with any set of desired functionalities. Further, we expect the advantage of the generative design to become more apparent when it is used to perform a targeted search in the latent space for polymers with a more complex set of properties.

## METHODS

### Unsupervised Learning: Polymer Variational Autoencoder.

A variational autoencoder<sup>50</sup> is a framework to learn the generative model with latent variables for the input dataset  $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N]$ . The generative process of each sample  $\mathbf{x}$  is performed by first sampling a latent vector  $\mathbf{z}$  from prior  $p(\mathbf{z})$  (typically a normal Gaussian distribution, denoted as  $\mathcal{N}(0, I)$ ) and then sample from conditional distribution  $p_\theta(\mathbf{x}|\mathbf{z})$  (known as “decoder”). The overall goal is to maximize the marginal likelihood  $p(\mathbf{x})$ , but due to the

```

smiles → chain | AT chain AT | AT chain
AT → ATT | ATT bond | bond ATT
ATT → '['*']
atom → bracket_atom | aliphatic_organic | aromatic_organic
aliphatic_organic → 'B' | 'C' | 'N' | 'O' | 'S' | 'P' | 'F' | 'T' | 'Cl' | 'Br' | 'H'
aromatic_organic → 'c' | 'n' | 'o' | 's'
bracket_atom → '[' BAI '['
BAI → isotope symbol BAC | symbol BAC | isotope symbol | symbol
BAC → chiral BAH | BAH | chiral
BAH → hcount charge | charge | hcount
symbol → aliphatic_organic | aromatic_organic
isotope → DIGIT | DIGIT DIGIT | DIGIT DIGIT DIGIT
DIGIT → '1' | '2' | '3' | '4' | '5' | '6' | '7' | '8' | '9'
chiral → '@' | '@@'
hcount → 'H' | 'H' DIGIT
charge → '-' | '-' DIGIT | '+' | '+' DIGIT
bond → '-' | '=' | '#' | '/' | '\'
ringbond → DIGIT | bond DIGIT
branched_atom → atom | atom BB | atom RB | atom RB BB
RB → ringbond | ringbond RB
BB → branch | branch BB
branch → '(' chain ')' | '(' bond chain ')' | '(' AT ')' | '(' chain AT ')' | '(' bond chain AT ')'
chain → branched_atom | branched_atom chain | branched_atom bond chain

```

The parse tree can be decomposed into a sequence of production rules and can be converted into a list of 1-hot vectors, with each dimension of the vector corresponding to a rule or a terminal symbol in the polymer SMILES grammar. Thus, each polymer SMILES gets converted into matrix  $\mathbf{X} \in \mathbb{R}^{T \times p}$ , where  $p$  denotes the total number of production rules and terminal symbols in the polymer grammar, and  $T$  is the number of productions applied to parse the input SMILES.

intractability of the direct optimization, a variational posterior  $p_\phi(\mathbf{z}|\mathbf{x})$  (known as “encoder”) is introduced to optimize the evidence lower bound (ELBO)

$$\begin{aligned} \mathcal{L}(\phi, \theta; \mathbf{x}) &= \log \int_{\mathbf{z}} p_\theta(\mathbf{x}|\mathbf{z})p(\mathbf{z}) \\ &\geq \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z}) + \log p(\mathbf{z}) - \log q_\phi(\mathbf{z}|\mathbf{x})] \end{aligned} \quad (1)$$

To make it easy to sample or calculate the probability from  $q_\phi(\mathbf{z}|\mathbf{x})$ , typically a parametric distribution is assumed (in this work, we assume a multivariate gaussian with a diagonal covariance matrix). By limiting the dimensionality of  $\mathbf{z}$ , the encoder is forced to learn meaningful representations in latent space. With a carefully designed decoder, the decoding probability mass will be reshaped toward the desired target region. Next, we introduce two important concepts of context-free-grammars and stochastically lazy attributes, which enforce the decoder to (re)construct syntactically and semantically valid polymer SMILES.

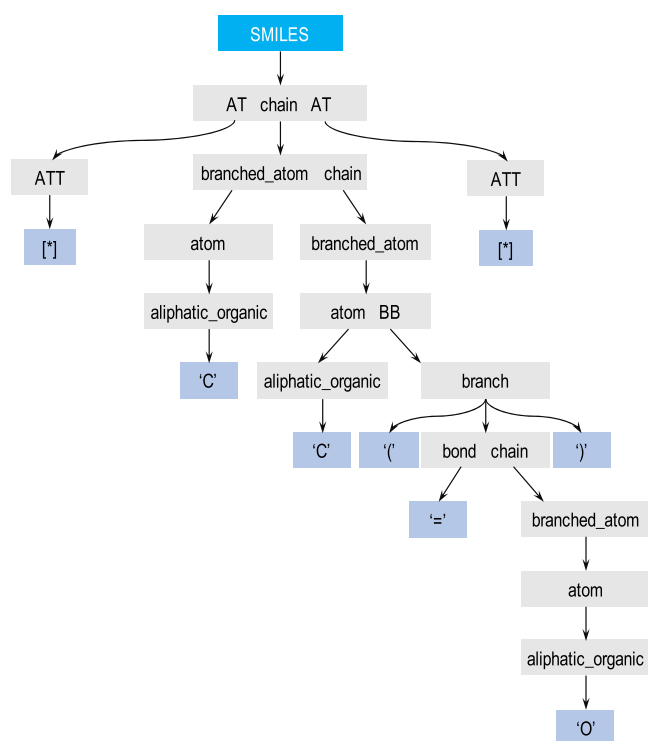
**Incorporating Grammar or Syntax.** Given a language and its context-free-grammar (CFG), a sentence in this language (e.g., a SMILES string) can be transformed to a parse tree.<sup>51</sup> While this concept is common for processing sentences in natural or programming languages, here we utilize it to transform polymer SMILES into its syntactic tree structure, which will be used as an input/output to the VAE. Formally, a CFG is defined as  $G = (V, \Sigma, R, S)$ , where  $V$  is a finite set of nonterminal symbols;  $\Sigma$  is a finite set of terminal symbols, distinct from  $V$ ;  $R$  is a finite set of production rules; and  $S \in V$  is a distinct nonterminal start symbol. Each production rule  $r \in R$  is denoted as  $r = \alpha \rightarrow \beta$  for  $\alpha \in V$  a nonterminal symbol and  $\beta = u_1, u_2, \dots, u_{|\beta|} \in (V \cup \Sigma)^*$  is a sequence of terminal and/or nonterminal symbols. Thus, for a string in a language—i.e., a sequence of terminals in  $\Sigma$ —the grammar  $G$  converts it into a syntactic tree, consisting of branching of nonterminal symbols in  $V$ , produced by recursively applying rules in  $R$  to leaf nodes, until all leaf nodes are terminal symbols in  $\Sigma$ . For example, a parse tree for a polyketone, with SMILES [\*]CC(=O)[\*], is shown in Figure 6; the “[\*]” symbolizes the chain ends. This was generated using the grammar  $G$  provided below:

For consistent dimensions across different SMILES,  $\mathbf{X}$  is padded with dummy numbers until  $T = T_{\max}$ . The matrix  $\mathbf{X}$  is then used as the input to the VAE encoder.

The VAE decoder attempts to reconstruct this matrix  $\mathbf{X}$  or the sequence of production rules (and thus, the polymer SMILES). To ensure syntactically valid SMILES are generated, at any time during the decoding process, the decoder is only allowed to select from a



subset of valid production rules. More details on the implementation of this process can be found here (Figure 6).<sup>26</sup>



**Figure 6.** Syntactic tree for polymer SMILES [\*]CC(=O)[\*] denoting a polyketone. The tree is generated by recursively applying production rules present in the grammar (see the text for details), until all leaf nodes contain terminal symbols.

**Incorporating Semantics.** The CFG implementation only guarantees production of syntactically valid SMILES. However, another crucial component for producing valid SMILES is to incorporate correct semantics or chemical rules. For example, the SMILES [\*]C, [\*]C(=O)C[\*], and [\*]c1ccccc[\*] are all syntactically correct but semantically invalid due to the absence of the second polymer chain end, unrealistic triple bond with O atom, and nonclosure of ringbond numbered 1, respectively.

In this work, we extend the previously proposed SD-VAE<sup>27</sup> to the polymer domain. The main idea of SD-VAE is to use the concept of attribute grammar to regulate the generation of the CFG tree by forcing the decoder to select from a smaller subset of semantically valid production rules. Formally, the attribute grammar is implemented by associating certain attribute(s) (or tags) to the nonterminal symbols of the CFG, written as  $\langle v \rangle.a$  for  $v \in V$ . Depending on the direction of information flow, attributes can either be inherited (parent-to-child nodes) or synthesized (child-to-parent nodes). However, these traditional concepts of attribute grammar can only be applied once the entire CFG tree of a molecule/polymer has been generated, making them inadequate for semantic checks. In contrast, SD-VAE proposed a generic way to introduce the concept of “lazy attribute” that allows performing the attribute grammar check alongside the generation of the CFG tree. More details on the implementation of SD-VAE for molecules are available in the original work.<sup>27</sup> Here, we extend it for the case of polymers by introducing a novel polymerization semantics constraint that guarantees the generated SMILES string has two chain ends with matched bond types. Furthermore, the framework extensions allow it to be applicable for both molecule and polymer discovery.

To realize the polymerization semantics, the following components are implemented in the generic SD-VAE framework:

- **Polymer inherited attribute:** this attribute is computed in the parent-to-child direction. If during polymer generation one chain end ([\*]) has already been generated with a particular bond type, say single bond, it will constraint the next generated chain end to also have the same bond type. Further, it is used to ensure that exactly two chain ends are generated.
- **Polymer synthesized attribute:** this attribute is computed in the child-to-parent direction. Once the first symbol [\*] and the corresponding bond type are generated during decoding, it is used to pass this information from the respective leaf node all of the way to the root node.
- **Polymer stochastic lazy attribute:** this attribute is exclusively introduced in the context of SD-VAE. It is used to memorize only the existence of [\*], while delaying the calculation of the synthesized attribute, i.e., the bond type associated with [\*]. Once the actual [\*] is generated, its bond type will be realized, and the synthesized attribute will be used to pass this information up the tree. Later, inherited attributes will be utilized for generation of the other matching [\*].

An example of attribute grammar for polymer tree generation is provided in the SI. Besides these polymer semantics, we also enforced semantics common to molecular systems. Thus, overall, the following set of semantics was imposed:

- **Polymer chain:** there should be exactly 2 [\*] to denote the chain ends; the bond type associated with the beginning and end [\*] should also match.
- **Valency constraints:** limit the maximum number of bonds associated with each atom.
- **Ringbond consistency:** ringbonds should come in pairs and could overlap with each other.
- **Aromatic:** atom type can only be aromatic when it is part of a ring.

**VAE Architecture.** For the encoder, we used 3-layer one-dimensional convolution neural networks (CNNs) followed by a fully connected dense layer, the output of which are used to get mean and variance of the distribution  $q_\phi(\mathbf{z}|\mathbf{x})$  through the reparametrization trick. To decode, 3-layer gated recurrent unit (GRU)-based recurrent neural network (RNN) is used followed by an affine layer activated by softmax to give probabilities  $p_\theta(\mathbf{x}|\mathbf{z})$ . The VAE architecture is inspired from the recent work on molecules,<sup>25–27</sup> allowing us to compare our results with past works. The value of  $T_{\max}$  was set to 512, while two VAE models with latent dimensionality of 64 and 128 were trained using Adam optimization algorithm.<sup>52</sup> The number of training epochs were determined by monitoring the model performance on a validation set of 20 000 SMILES taken from the hypothetical SMILES dataset. The VAE code was implemented in the PyTorch framework.

#### Supervised Learning: Gaussian Process Regression (GPR).

After training the polymer VAE, we use GPR to learn polymer properties using their encoded latent representations. GPR uses a Bayesian framework, wherein a Gaussian process is used to obtain the functional mapping  $f(\mathbf{z}) \rightarrow y$  based on the available training set and the Bayesian prior, incorporated using the covariance (or kernel) function. The square exponential kernel,  $k(\mathbf{z}, \mathbf{z}') = \sigma_f \exp\left(-\frac{1}{2\sigma_l^2} \|\mathbf{z} - \mathbf{z}'\|^2\right) + \delta_{ij}\sigma_n^2$ , with three hyperparameters  $\sigma_f$ ,  $\sigma_l$ , and  $\sigma_n$  was chosen for this work. This choice of kernel is quite standard and is known to work well for a variety of materials problems.

**Training Datasets.** Deep generative models, such as VAE, need larger datasets (in the order of 100 000 examples) for the training to be effective. While such large datasets are available for the case of molecules, we could only obtain information on  $\sim 12$  000 polymers, experimentally reported in the literature. Thus, using this initial dataset, a strategy was devised to create a larger hypothetical polymer SMILES dataset. The strategy is based on the use of retro-synthesis principles to find various polymer “building blocks”, which are then combined together to form hypothetical polymers. Care is taken to preserve the frequency of occurrence of different building blocks in the hypothetical dataset, making it realistic and representative of the

initial empirically known polymers. A total of ~250 000 hypothetical polymer SMILES were created by combining the various number of building blocks, ranging from 2 to 7 (see the SI). Except 20 000 SMILES used as the validation set, all of the remaining hypothetical polymer SMILES were used for VAE training. Model parameters with the lowest validation error were selected.

GPR property prediction models were trained using empirically known  $T_g$  and theoretically computed  $E_g$  datasets containing 4997 and 963 data points, respectively. While  $T_g$  dataset was collected from experimental reports in the literature, the  $E_g$  dataset consisted of density functional theory (DFT)-computed bandgap values of polymer single chains. It should be noted that measuring  $E_g$  values for polymers is a nontrivial task, and the DFT-computed values can only be considered as reasonable estimates. Nonetheless, past works suggest a good correlation between the DFT-computed and the empirically measured  $E_g$  values for a few common polymers.<sup>53</sup>

## ■ ASSOCIATED CONTENT

### SI Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.chemmater.0c03332>.

List of the new polymers (SMILES and VAE property predictions) designed in this work, and additional discussion on the VAE-128 latent space, polymer bandgap validation, hypothetical SMILES dataset, chemical fragment analysis, and polymer attribute grammar (PDF)

Publication\_data\_v3 (ZIP)

## ■ AUTHOR INFORMATION

### Corresponding Author

**Rohit Batra** – School of Materials Science & Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332, United States; Center for Nanoscale Materials, Argonne National Laboratory, Lemont, Illinois 60439, United States; [orcid.org/0000-0002-1098-7035](https://orcid.org/0000-0002-1098-7035); Email: [rohitbatra1989@gmail.com](mailto:rohitbatra1989@gmail.com)

### Authors

**Hanjun Dai** – Computational Science & Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332, United States

**Tran Doan Huan** – School of Materials Science & Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332, United States

**Lihua Chen** – School of Materials Science & Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332, United States; [orcid.org/0000-0002-9852-8211](https://orcid.org/0000-0002-9852-8211)

**Chiho Kim** – School of Materials Science & Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332, United States; [orcid.org/0000-0002-1814-4980](https://orcid.org/0000-0002-1814-4980)

**Will R. Gutekunst** – School of Chemistry & Biochemistry, Georgia Institute of Technology, Atlanta, Georgia 30332, United States

**Le Song** – Computational Science & Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332, United States

**Rampi Ramprasad** – School of Materials Science & Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332, United States; [orcid.org/0000-0003-4630-1565](https://orcid.org/0000-0003-4630-1565)

Complete contact information is available at: <https://pubs.acs.org/doi/10.1021/acs.chemmater.0c03332>

## Notes

The authors declare no competing financial interest.

## ■ ACKNOWLEDGMENTS

This work is supported by the Office of Naval Research through N0014-17-1-2656, a Multi-University Research Initiative (MURI) grant. R.B. acknowledges the use of the Center for Nanoscale Materials, an Office of Science user facility, and support by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under Contract No. DE-AC02-06CH11357. The use of the Carbon computing cluster at the Center for Nanoscale materials at the Argonne National Laboratory is gratefully acknowledged. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

## ■ REFERENCES

- (1) Meredig, B.; Agrawal, A.; Kirklín, S.; Saal, J. E.; Doak, J.; Thompson, A.; Zhang, K.; Choudhary, A.; Wolverton, C. Combinatorial screening for new materials in unconstrained composition space with machine learning. *Phys. Rev. B* **2014**, *89*, No. 094104.
- (2) Pilia, G.; McClellan, K. J.; Stanek, C. R.; Uberuaga, B. P. Physics-informed machine learning for inorganic scintillator discovery. *J. Chem. Phys.* **2018**, *148*, No. 241729.
- (3) Chen, L.; Tran, H.; Batra, R.; Kim, C.; Ramprasad, R. Machine learning models for the lattice thermal conductivity prediction of inorganic materials. *Comput. Mater. Sci.* **2019**, *170*, No. 109155.
- (4) Tshitoyan, V.; Dagdelen, J.; Weston, L.; Dunn, A.; Rong, Z.; Kononova, O.; Persson, K. A.; Ceder, G.; Jain, A. Unsupervised word embeddings capture latent knowledge from materials science literature. *Nature* **2019**, *571*, 95.
- (5) Sun, S.; Ouyang, R.; Zhang, B.; Zhang, T.-Y. Data-driven discovery of formulas by symbolic regression. *MRS Bull.* **2019**, *44*, 559–564.
- (6) Vasudevan, R. K.; Choudhary, K.; Mehta, A.; Smith, R.; Kusne, G.; Tavazza, F.; Vlcek, L.; Ziatdinov, M.; Kalinin, S. V.; Hatrick-Simpers, J. Materials science in the artificial intelligence age: high-throughput library generation, machine learning, and a pathway from correlations to the underpinning physics. *MRS Commun.* **2019**, 821–838.
- (7) Kim, E.; Huang, K.; Saunders, A.; McCallum, A.; Ceder, G.; Olivetti, E. Materials synthesis insights from scientific literature via text extraction and machine learning. *Chem. Mater.* **2017**, *29*, 9436–9444.
- (8) King, R. D.; Rowland, J.; Oliver, S. G.; Young, M.; Aubrey, W.; Byrne, E.; Liakata, M.; Markham, M.; Pir, P.; Soldatova, L. N.; Sparkes, A.; Whelan, K. E.; Clare, A. The automation of science. *Science* **2009**, *324*, 85–89.
- (9) Kim, C.; Chandrasekaran, A.; Huan, T. D.; Das, D.; Ramprasad, R. Polymer genome: a data-powered polymer informatics platform for property predictions. *J. Phys. Chem. C* **2018**, *122*, 17575–17585.
- (10) Xue, D.; Balachandran, P. V.; Hogden, J.; Theiler, J.; Xue, D.; Lookman, T. Accelerated search for materials with targeted properties by adaptive design. *Nat. Commun.* **2016**, *7*, No. 11241.
- (11) Mansouri Tehrani, A.; Oliynyk, A. O.; Parry, M.; Rizvi, Z.; Couper, S.; Lin, F.; Miyagi, L.; Sparks, T. D.; Brgoch, J. Machine learning directed search for ultraincompressible, superhard materials. *J. Am. Chem. Soc.* **2018**, *140*, 9844–9853.
- (12) Ren, F.; Ward, L.; Williams, T.; Laws, K. J.; Wolverton, C.; Hatrick-Simpers, J.; Mehta, A. Accelerated discovery of metallic glasses through iteration of machine learning and high-throughput experiments. *Sci. Adv.* **2018**, *4*, No. eaq1566.

- (13) Stanev, V.; Oses, C.; Kusne, A. G.; Rodriguez, E.; Paglione, J.; Curtarolo, S.; Takeuchi, I. Machine learning modeling of superconducting critical temperature. *npj Comput. Mater.* **2018**, *4*, No. 29.
- (14) Lopez, S. A.; Sanchez-Lengeling, B.; de Goes Soares, J.; Aspuru-Guzik, A. Design principles and top non-fullerene acceptor candidates for organic photovoltaics. *Joule* **2017**, *1*, 857–870.
- (15) Kanal, I. Y.; Owens, S. G.; Bechtel, J. S.; Hutchison, G. R. Efficient computational screening of organic polymer photovoltaics. *J. Phys. Chem. Lett.* **2013**, *4*, 1613–1623.
- (16) Cheng, L.; Assary, R. S.; Qu, X.; Jain, A.; Ong, S. P.; Rajput, N. N.; Persson, K.; Curtiss, L. A. Accelerating electrolyte discovery for energy storage with high-throughput screening. *J. Phys. Chem. Lett.* **2015**, *6*, 283–291.
- (17) Chandrasekaran, A.; Kamal, D.; Batra, R.; Kim, C.; Chen, L.; Ramprasad, R. Solving the electronic structure problem with machine learning. *npj Comput. Mater.* **2019**, *5*, No. 22.
- (18) Huan, T. D.; Batra, R.; Chapman, J.; Krishnan, S.; Chen, L.; Ramprasad, R. A universal strategy for the creation of machine learning-based atomistic force fields. *npj Comput. Mater.* **2017**, *3*, No. 37.
- (19) Pun, G. P.; Batra, R.; Ramprasad, R.; Mishin, Y. Physically informed artificial neural networks for atomistic modeling of materials. *Nat. Commun.* **2019**, *10*, No. 2339.
- (20) Huan, T. D.; Batra, R.; Chapman, J.; Kim, C.; Chandrasekaran, A.; Ramprasad, R. Iterative-learning strategy for the development of application-specific atomistic force fields. *J. Phys. Chem. C* **2019**, *123*, 20715–20722.
- (21) Shahriari, B.; Swersky, K.; Wang, Z.; Adams, R. P.; de Freitas, N. Taking the human out of the loop: A review of bayesian optimization. *Proc. IEEE* **2016**, *104*, 148–175.
- (22) Kim, C.; Chandrasekaran, A.; Jha, A.; Ramprasad, R. Active-learning and materials design: the example of high glass transition temperature polymers. *MRS Commun.* **2019**, 860–866.
- (23) Walsh, A. Inorganic materials: The quest for new functionality. *Nat. Chem.* **2015**, *7*, 274.
- (24) Kim, C.; Batra, R.; Chen, L.; Tran, H.; Ramprasad, R. Polymer design using genetic algorithm and machine learning. *Comput. Mater. Sci.* **2021**, *186*, No. 110067.
- (25) 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 Cent. Sci.* **2018**, *4*, 268–276.
- (26) Kusner, M. J.; Paige, B.; Hernández-Lobato, J. M. In *Grammar Variational Autoencoder*, Proceedings of the 34th International Conference on Machine Learning-Volume 70, 2017; pp 1945–1954.
- (27) Dai, H.; Tian, Y.; Dai, B.; Skiena, S.; Song, L. Syntax-Directed Variational Autoencoder for Structured Data. 2018, arXiv:1802.08786. arXiv.org e-Print archive. <https://arxiv.org/abs/1802.08786> (accessed Nov 21, 2020).
- (28) Jin, W.; Barzilay, R.; Jaakkola, T. Junction Tree Variational Autoencoder for Molecular Graph Generation. 2018, arXiv:1802.04364, arXiv.org e-Print archive. <https://arxiv.org/abs/1802.04364> (accessed Nov 21, 2020).
- (29) Liu, Q.; Allamanis, M.; Brockschmidt, M.; Gaunt, A. Constrained graph variational autoencoders for molecule design. *Adv. Neural Inf. Process. Syst.* **2018**, 7795–7804.
- (30) 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. 2017, arXiv:1705.10843, arXiv.org e-Print archive. <https://arxiv.org/abs/1705.10843> (accessed Nov 21, 2020).
- (31) You, J.; Liu, B.; Ying, Z.; Pande, V.; Leskovec, J. Graph convolutional policy network for goal-directed molecular graph generation. *Adv. Neural Inf. Process. Syst.* **2018**, 6410–6421.
- (32) De Cao, N.; Kipf, T. MolGAN: An Implicit Generative Model for Small Molecular Graphs. 2018, arXiv:1805.11973, arXiv.org e-Print archive, <https://arxiv.org/abs/1805.11973> (accessed Nov 21, 2020).
- (33) 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. *Nat. Mach. Intell.* **2020**, *2*, 254–265.
- (34) Krenn, M.; Häse, F.; Nigam, A.; Friederich, P.; Aspuru-Guzik, A. Self-referencing embedded strings (SELFIES): A 100% robust molecular string representation. *Mach. Learn.: Sci. Technol.* **2020**, *1*, No. 045024.
- (35) Jørgensen, P. B.; Mesta, M.; Shil, S.; García Lastra, J. M.; Jacobsen, K. W.; Thygesen, K. S.; Schmidt, M. N. Machine learning-based screening of complex molecules for polymer solar cells. *J. Chem. Phys.* **2018**, *148*, No. 241735.
- (36) Kim, B.; Lee, S.; Kim, J. Inverse design of porous materials using artificial neural networks. *Sci. Adv.* **2020**, *6*, No. eaax9324.
- (37) Allamanis, M.; Chanthirasegaran, P.; Kohli, P.; Sutton, C. In *Learning Continuous Semantic Representations of Symbolic Expressions*, Proceedings of the 34th International Conference on Machine Learning-Volume 70, 2017; pp 80–88.
- (38) Kernighan, B. W.; Ritchie, D. M.; Ekelint, P. *The C Programming Language*; Prentice-Hall Englewood Cliffs, 1988; Vol. 2.
- (39) St. John, P. C.; Phillips, C.; Kemper, T. W.; Wilson, A. N.; Guan, Y.; Crowley, M. F.; Nimlos, M. R.; Larsen, R. E. Message-passing neural networks for high-throughput polymer screening. *J. Chem. Phys.* **2019**, *150*, No. 234111.
- (40) Jin, W.; Barzilay, R.; Jaakkola, T. Hierarchical Generation of Molecular Graphs Using Structural Motifs. 2020, arXiv:2002.03230, arXiv.org e-Print archive. <https://arxiv.org/abs/2002.03230> (accessed Nov 21, 2020).
- (41) Ho, J. S.; Greenbaum, S. G. Polymer capacitor dielectrics for high temperature applications. *ACS Appl. Mater. Interfaces* **2018**, *10*, 29189–29218.
- (42) Tan, D.; Zhang, L.; Chen, Q.; Irwin, P. High-temperature capacitor polymer films. *J. Electron. Mater.* **2014**, *43*, 4569–4575.
- (43) Zhou, Y.; Li, Q.; Dang, B.; Yang, Y.; Shao, T.; Li, H.; Hu, J.; Zeng, R.; He, J.; Wang, Q. A scalable, high-throughput, and environmentally benign approach to polymer dielectrics exhibiting significantly improved capacitive performance at high temperatures. *Adv. Mater.* **2018**, *30*, No. 1805672.
- (44) Wu, C.; Deshmukh, A. A.; Li, Z.; Chen, L.; Alamri, A.; Wang, Y.; Ramprasad, R.; Sotzing, G. A.; Cao, Y. Flexible temperature-invariant polymer dielectrics with large bandgap. *Adv. Mater.* **2020**, *32*, No. 2000499.
- (45) Mannodi-Kanakithodi, A.; Chandrasekaran, A.; Kim, C.; Huan, T. D.; Pilania, G.; Botu, V.; Ramprasad, R. Scoping the polymer genome: A roadmap for rational polymer dielectrics design and beyond. *Mater. Today* **2018**, *21*, 785–796.
- (46) Kim, C.; Pilania, G.; Ramprasad, R. From organized high-throughput data to phenomenological theory using machine learning: the example of dielectric breakdown. *Chem. Mater.* **2016**, *28*, 1304–1311.
- (47) Kamal, D.; Wang, Y.; Tran, H. D.; Chen, L.; Li, Z.; Wu, C.; Nasreen, S.; Cao, Y.; Ramprasad, R. Computable bulk and interfacial electronic structure features as proxies for dielectric breakdown of polymers. *ACS Appl. Mater. Interfaces* **2020**, 37182–37187.
- (48) Lin, T.-S.; Coley, C. W.; Mochigase, H.; Beech, H. K.; Wang, W.; Zhang, Z.; Woods, E.; Craig, S. L.; Johnson, J. A.; Kalow, J. A.; et al. BigSMILES: a structurally-based line notation for describing macromolecules. *ACS Cent. Sci.* **2019**, *5*, 1523–1531.
- (49) Zhou, Z.; Kearnes, S.; Li, L.; Zare, R. N.; Riley, P. Optimization of molecules via deep reinforcement learning. *Sci. Rep.* **2019**, *9*, No. 10752.
- (50) Kingma, D. P.; Welling, M. In *Auto-Encoding Variational Bayes*, Proceedings of the International Conference on Learning Representations (ICLR), 2014.
- (51) Hopcroft, J. E.; Motwani, R.; Ullman, J. D. *Introduction to Automata Theory, Languages and Computation*; Addison-Wesley, 2006.
- (52) Kingma, D. P.; Ba, J. In *Adam: A Method for Stochastic Optimization*, International Conference on Learning Representations, 2015.

(53) Chen, L.; Venkatram, S.; Kim, C.; Batra, R.; Chandrasekaran, A.; Ramprasad, R. Electrochemical stability window of polymeric electrolytes. *Chem. Mater.* **2019**, *31*, 4598–4604.