THE JOURNAL OF PHYSICAL CHEMISTRY A
A JOURNAL OF THE AMERICAN CHEMICAL SOCIETY

pubs.acs.org/JPCA

Article

Actively Searching: Inverse Design of Novel Molecules with Simultaneously Optimized Properties

Nicolae C. Iovanac, Robert MacKnight, and Brett M. Savoie*

Cite This: J. Phys. Chem. A 2022, 126, 333−340

Read Online

ACCESS | Metrics & More | Article Recommendations | Supporting Information

ABSTRACT: Combining quantum chemistry characterizations with generative machine learning models has the potential to accelerate molecular discovery. In this paradigm, quantum chemistry acts as a relatively cost-effective oracle for evaluating the properties of particular molecules, while generative models provide a means of sampling chemical space based on learned structure−function relationships. For practical applications, multiple potentially orthogonal properties must be optimized in tandem during a discovery workflow. This carries additional difficulties associated with the specificity of the targets and the ability for the model to reconcile all properties simultaneously. Here, we demonstrate an active learning approach to improve the performance of multi-target generative chemical models. We first demonstrate the effectiveness of a set of baseline models trained on single property prediction tasks in generating novel compounds (i.e., not present in the training data) with various property targets, including both interpolative and extrapolative generation scenarios. For property ranges where accurate targeting proves difficult, the novel compounds suggested by the model are characterized using quantum chemistry and the new molecules closest to expressing the desired properties are fed back into the generative model for additional training. This gradually improves the generative models' understanding of targeted areas of chemical space and shifts the distribution of the generated compounds toward the targeted values. We then demonstrate the effectiveness of this active learning approach in generating compounds with multiple chemical constraints, including vertical ionization potential, electron affinity, and dipole moment targets, and validate the results at the ωB97X-D3/def2-TZVP level. This method requires no modifications to extant generative approaches, but rather utilizes their inherent generative and predictive aspects for self-refinement, and can be applied to situations where any number of properties with varying degrees of correlation must be optimized simultaneously.

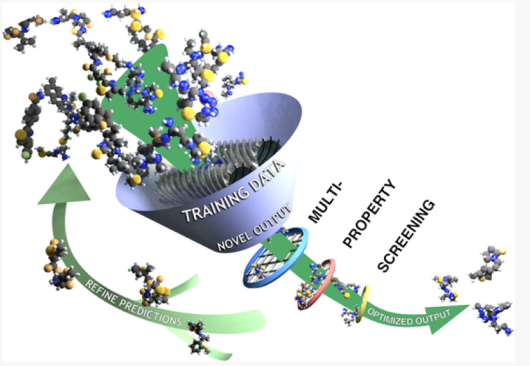

INTRODUCTION

Machine learning (ML) has emerged as a powerful tool for solving previously intractable problems by extracting latent information from domain data and has been effectively employed in areas as distinct as manufacturing analytics[1] and cancer detection.[2] In recent years, ML has been proved particularly successful in the chemical sciences, where it is used to predict interatomic potentials,[3] quantum chemical properties,[4] and structural data of polymers[5,6] and crystals.[7] Moving beyond the "forward-problem" of predicting molecular properties from a given chemical structure, generative chemical models have garnered significant interest in solving the "inverse-problem" of predicting a chemical structure from a given descriptor. As a large body of research in chemistry is devoted to creating novel compounds under functional constraints, these generative models have the potential to supplement and automate much of the often-laborious manual optimization by providing reasonable chemical suggestions for more expensive experimental synthesis and characterization. Generative adversarial networks[8−11] and various formulations of autoencoder networks[12−18] have emerged as some of the more popular frameworks for a generative machine-learning-based chemical design. These methods often accomodate property prediction that allows suggestions to be biased toward compounds with particular properties.[19] Much effort has been directed to solving issues related to the ability of these models to generate valid chemistries,[9,20−22] and they have been successfully demonstrated in generating compounds with specific properties such as the band gap[23] and thermal conductivity.[24]

While models capable of optimizing one molecular property are compelling proof-of-principle demonstrations, multi-property optimization is required in any practical chemical discovery application. Because of the exponential scaling of search spaces with respect to the number of properties, multi-property chemical searches are fundamentally more challeng-

Received: September 16, 2021
Revised: December 16, 2021
Published: January 5, 2022

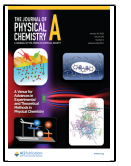

https://doi.org/10.1021/acs.jpca.1c08191
J. Phys. Chem. A 2022, 126, 333−340

ing because they will be typically operating in an extrapolative regime (i.e., searching for properties outside the convex hull of training data ranges) and training data density drops in high dimensions. Several recent studies have highlighted the challenges and potential solutions to pursuing multi-property searches. Janet et al. balanced the solubility and redox potential in the design of transition-metal complexes for redox flow batteries using global optimization to explore an enumerated space of 2.8 million candidate complexes.[25] Domenico et al. utilized reinforcement learning for the design of drug-like molecules where the trade-offs among relevant physiochemical properties like molecular weight and hydrogen bond donors/acceptors, as well as similarity constraints to known drugs, were minimized.[26] Ståhl et al. also used a reinforcement learning approach to target and modify fragments in known structures to develop novel structures similar to known lead compounds but with an optimized molecular weight, log *P*, and polar surface area.[27] Nigam et al. recently proposed the STONED algorithm, which side-steps the data limitations associated with training deep generative models and instead relies on string permutations of seed structures represented with semantically robust SELFIES.[28] Zhou et al. developed a reinforcement learning method based on atom/bond addition/removal to optimize compounds with respect to log *P* and a quantitative estimate of drug likeliness (QED).[29] Interestingly, they also note that common targets for generative models may not be suitable for real world applications. Log *P*, for instance, may be trivially improved by simply increasing the length of carbon chains in a structure. Of interest is a method that can be applied generally to experimental properties or computational analogues.

Despite the large data sets available (and in many cases necessary) for training, certain combinations of properties are difficult for a generative model to achieve, either because they contradict basic physical relationships, or because they simply have a limited representation within the training data. Rather than filtering an enumerated set of compounds or guiding the generation process with methods such as reinforcement learning, we propose leveraging the generative aspect of these models to enrich training data in the targeted regions of a chemical space. Generative chemical models have the unique feature that syntactically valid outputs are guaranteed to belong within the chemical space, meaning that they are suitable samples for further model training. By sampling under-represented regions of chemical space, new compounds may be discovered that are closer to the desired property space than any elements in the training set. By introducing the model to these new chemistries, the model can better learn which features correlate with the designated figures of merit. This framework falls under the paradigm of active learning. In active learning, a model can ask an expert source (i.e., the oracle) to annotate unlabeled training data that the model believes will be helpful.[30] This is particularly useful in situations where generating labeled data is difficult, as is often the case with chemical property data, because in an optimal case the model will utilize as little data as possible. This method was exploited by Konze et al. who used an active learning-based approach to more efficiently screen a large set of ligands without conducting expensive free-energy perturbations on the entire set.[31] This approach was extended in their recent follow-up to train a goal-directed generative model to generate promising ligands for further screening.[32] The approach of selecting training samples based on proximity to an optimization objective is well represented in materials design applications, such as a search for polymers with high glass transition temperatures[33] or CO adsorption sites in electrocatalysts.[34] Other active learning strategies also include prediction uncertainty as part of their acquisition function for prioritizing samples for characterization. The cost of data generation has led to a proliferation of similar approaches to efficiently select optimal training data, all falling under the active learning designation.

We propose formulating the entire goal of multi-target chemical optimization as an active learning problem. Rather than attempting to determine optimal regions of the chemical space to sample or train on, we query the model to obtain its suggestions for compounds with the desired properties. These compounds are then screened via a semiempirical calculation (i.e., the oracle), and the model is retrained on those new compounds that actually match the target property profiles. In this way, the model develops its own training data to better understand new chemistries, and iterating on this procedure provides an opportunity to continuously improve a model's ability to target compounds with under-represented properties.

Herein, we examine this active learning framework to improve the performance of multi-objective generative chemical models. Utilizing a subset of compounds from the ZINC15 database,[35] we develop our own data set of quantum chemical properties including the vertical ionization potential (VIP), electron affinity (EA), and dipole moment (DM) calculated at the semiempirical level for training. All three properties are relatively unexplored in generative applications and are essential characteristics in any organic electronics application. In particular, VIP and EA determine redox behaviors, while DM is an important contribution to the bulk dielectric and impacts the creation of charge traps in charge transport applications. Likewise, we wanted to investigate properties besides the quantitative estimate of drug likeness (QED) and the water−octanol partition coefficient (log *P*), which are over-represented in the generative literature due to their ease of calculation and data abundance. We demonstrate the effectiveness of generative chemical models trained to propose compounds with a single targeted property, as well as multiple properties at once. We also demonstrate the shortcomings of such models, particularly when the combination of desired properties is not found in the training data, and how they may be overcome with active learning. We then validate the properties of the newly suggested compounds at the ωB97X-D3/def2-TZVP level. This active learning scheme can be applied to both single property and multi-property models to extrapolate to new regions of chemical space.

METHODOLOGY

Data Sets. All models were initially trained and evaluated using structures from the ZINC15 data set. This data set contains 3D structural data for hundreds of millions of small molecules, from which we chose a subset of 250,000 compounds with a molecular weight between 200 and 500 Da and log *P* between −1 and 5. These compounds were subjected to geometry optimization and electronic structure calculation with GFN2-xTB[36] (xTB) to obtain their DM, VIP, and EA. After removing compounds that failed the initial geometry optimization, we were left with 224,742 structures and their associated properties. Eighty percent were utilized for training, with the remaining 20 percent being withheld for

https://doi.org/10.1021/acs.jpca.1c08191
J. Phys. Chem. A 2022, 126, 333−340

validation. Additional structures generated during the active learning step (see below) were subjected to the same property calculation methods to expand the data set. The training data and xTB calculated properties for all the sampled structures from all of the generative active learning have been uploaded to Zenodo with a persistent DOI (10.5281/zenodo.5512896).

In order to validate the properties of the structures generated in the multi-objective active learning study, the xTB-optimized geometries were used as inputs for density functional theory (DFT) geometry optimizations at the ωB97X-D3/def2-TZVP level to determine the DMs and single-point energies of the neutral species.[37] Using the optimized neutral geometries, single-point calculations were also performed for the cation and anion electronic states to calculate the VIP and EA, respectively. All DFT calculations were conducted using Orca 4.0.1.[38]

All properties were calculated based on a single conformer that was initialized with all-trans relationships for the dihedrals involving heavy atoms, followed by force-field geometry optimization, then xTB optimization. It is significant that DM, VIP, and EA can all be affected by conformation, whereas the generative model used here is based on a graphical representation that does not include geometric information. Using a graphical representation amounts to coarse graining, where the model is tasked with learning property trends conditioned on the procedure used for generating the conformers. In the current case, we have applied a systematic procedure for generating the conformer of the structures, but this is nevertheless a potentially important source of error for any property prediction model based on a graphical featurization.

Machine Learning Architecture. Three models were developed for single target predictions and one model was developed for multi-property prediction tasks. We utilized the grammar variational autoencoder[39] (GVAE) model architecture to achieve the generation of molecules with the targeted properties. The GVAE architecture uses grammar parse trees to represent molecules, which are decoded to SMILES strings with rules to reduce syntactically invalid outputs and increase the reconstruction rate. We modified the baseline GVAE approach by utilizing a single linear predictor layer for each property being predicted.[40] A simple predictor has limited capability to reduce prediction loss, and instead the encoder must adjust its encodings to ensure that properties vary linearly within the latent space. In particular, we have previously shown how this architecture leads to effective transfer learning between the correlated property prediction tasks in data scarce scenarios.[40] Additionally, the linear prediction tasks force the encoder to map structures such that the principal components of the latent encodings now have a direct, linear relationship with the properties that the predictor is trained on. This allows for a simple procedure to sample new candidate molecules, as discussed in the next section. The implemented autoencoder accepts one-hot inputs of grammar parse trees to an encoder network comprising three one-dimensional convolutional layers with filter sizes of 9, 9, and 10 and kernel sizes of 9, 9, and 11. The outputs from the convolutional layers are passed to a fully connected layer of 435 units, which are then separately connected to two fully connected layers of 56 units (i.e., the dimensionality of the latent space) defining the mean and log variance of the encoding distribution. The decoder accepts samples from the encoding distributions and passes them to a fully connected layer of 56 units that is connected to three gated recurrent units of 501 cells each before terminating in a final fully connected layer that outputs the probability distributions for the output sequence. The latent space dimensionality determines the level of compression performed by the autoencoder and cannot be set too small without incurring a loss of accuracy in the property prediction tasks. A 56-dimensional latent space matches the original implementation by Kusner et al.[39] and was adopted here for its sufficient property prediction accuracy for the generated species. The effect of this and other hyperparameters were investigated in our earlier study for the interested reader.[40] ReLU activation functions were used for all units in the autoencoder. Property prediction is performed by passing latent vectors to a single linear unit producing a scalar output. An additional unit is included for each property. Models were created using Keras[41] 2.2.4 with Tensorflow[42] 1.14.0 backend. A diagram summarizing the model architecture has been included in Figure S1 of the Supporting Information.

It is difficult to train a network on both encoding/decoding and property prediction from scratch. It often proves much easier to first train the network on the encoding/decoding task first to obtain useful compressed chemical representations before transferring those weights for fine-tuning in a joint training task. For this purpose, the SMILES strings corresponding to all training compounds in the original GVAE data set were first converted into one-hot grammar parse trees and used as both inputs and labeled outputs for autoencoder pretraining. Pretraining followed the same routine used in the original GVAE implementation.[39] The pretrained model was then utilized for the weight initialization of the autoencoder for joint training on property prediction tasks, where in addition to encoding and decoding, the model was tasked with predicting up to three chemical properties from the latent encodings. Training was conducted using the RMSprop algorithm with a learning rate of 0.001, which was set to decay by a factor of 0.3 in the case of a plateau in the validation loss. In order to balance the performance across all tasks and to ensure stable training, the loss weights assigned to the different tasks were adjusted. Variational loss was scaled by 750, and the categorical cross-entropy loss from encoding/decoding was set to 50 initially, before decaying to 1 according to a sigmoid function. The MSE losses for the property predictors were not scaled, but all properties were normalized to fall in the range −20 to 20.

Sampling Paradigms. With a fully trained autoencoder, new molecules may be decoded from arbitrary points in the chemical latent space. Jointly training the autoencoder with a property prediction task based on a linear prediction network ensures that the property varies linearly along the first principal component of the latent encodings.[40] The use of a linear predictor results in a simple, interpretable latent space and thus obviates the need for a surrogate model to correlate latent position with properties of interest; compounds with specific properties can be generated by targeting regions of the latent space based on univariate linear regression between the property of interest and the position along the first principal component. In the case of multi-property models, each property will tend to vary linearly along a corresponding principal component. Linear regression along each of these axes determines the components of the sampled vector that would correspond to a compound displaying those properties; however, a correlation between the properties may lead to latent space organization not being exactly orthogonal. To find

https://doi.org/10.1021/acs.jpca.1c08191
J. Phys. Chem. A 2022, 126, 333−340

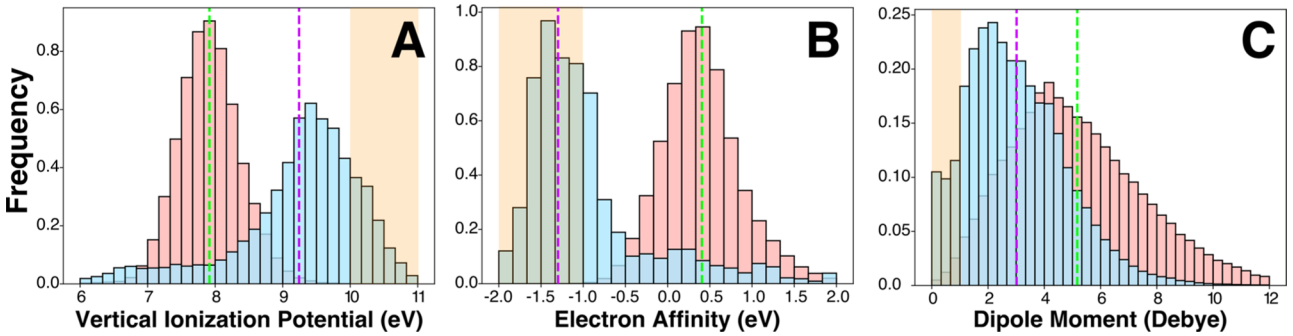

Figure 1. Property histograms for molecules generated by models trained on (A) VIP, (B) EA, and (C) DM. For each model, 100,000 structures were generated and subsequently characterized at the xTB level. Training distributions are shown in red with generated data in blue. Means of the training distributions and generated distributions are indicated by green- and purple-dotted lines, respectively. Models are tasked with generating compounds with properties ranges, shown in orange, that extrapolate beyond the values represented in the training data.

the directions to the sample, we linearly regress the angles that the points must be rotated by to maximize the R^2 between the position along a particular principal component and the property of interest. When optimizing multiple properties, the probabilities are jointly sampled to account for potential correlations between the latent space directions of the properties. All other latent dimensions are sampled normally with the mean and standard deviation determined by the training data. While this sampling approach only leverages linear correlations between properties, we note that other approaches involving surrogate models, like Gaussian processes regression, can be utilized for the sampling step.[19,43,44] These may be useful in situations where a linear predictor does not provide sufficient prediction accuracy or where the uncertainty-based sampling is desired.

Active Learning Technique. Depending on the property, the range for which it is sampled, and whether the search is extrapolative, the model may not accurately generate structures with properties that match those suggested by the regression outlined above. Nevertheless, structures sampled from the regions of the latent space that the model predicts will yield the targeted ranges are still potentially valuable for refining structure-property relationships in the targeted region. The molecules sampled from the targeted region were characterized with xTB and used for retraining the model as follows. In each iteration, 100,000 unique structures were sampled from the targeted property region, as predicted by the model, screened for validity, and the canonical SMILES of the valid structures were checked against the training and validation data sets to ensure the novelty of the generated structures. Novel structures were subjected to xTB calculations for characterization. Because the property prediction model is imperfect, these compounds were filtered after the xTB characterizations to ensure their properties were within the desired range before use in retraining. In situations where no compounds displayed properties within the desired range, we instead selected molecules with properties that were above or below (depending on the extrapolated target region) the median value in the training set, thus still providing the model with compounds exhibiting properties that were closer to the targeted range than the original training data. It was found that retraining the model solely on this newly generated data significantly harmed performance; this effect could be due to the significant differences between the new data and the original training data leading to catastrophic forgetting. Instead of only retraining with the new data, which due to the screening process always contains less than 100,000 total structures, compounds from the training set were randomly sampled and added to this new data set until reaching 100,000 total training structures. Our rationale for mixing in a fraction of the original training data is to give the model examples of structures that do not fulfill the targeted criteria while maintaining a training data size that trains effectively with fixed model hyperparameters. Training on all structures from all iterations is also a possible option but was not explored here. Because each iteration is performed on the same total number of structures, computational cost tends to scale linearly with the number of iterations performed. Depending on the specific number of properties being optimized and the relationship between the initial training data and the targeted property ranges, the data set size sampled at each iteration might need to be modified to ensure effective training and convergence. A minimum of five active learning iterations were specified for each experiment. For the single-property case, iterations were stopped after doubling the percentage of generated structures within the target region. For the multi-property case, iterations were stopped after obtaining 1500 molecules exhibiting the targeted property ranges at the xTB level to provide a tractable number for further DFT validation. These convergence criteria were selected out of convenience; alternative stopping criteria could be based on computational time, total number of iterations, or total number of generated structures in addition to proximity to property targets. All candidate rankings and subsequent trainings utilize the xTB-computed properties rather than those predicted by the model to provide external validation.

RESULTS AND DISCUSSION

Single Property Searches. In previous work, we have demonstrated that the use of chemical autoencoders for targeted structure searching is effective for properties within the GDB19 data set,[45] namely, internal energy, zero-point vibrational energy, and HOMO−LUMO gap.[46] To establish a baseline performance on the new properties without active learning, we investigated three models trained to individually predict VIP, EA, and DM, and sampled 100,000 structures in property ranges poorly represented in the training data. The targeted ranges for VIP, EA, and DM were 10.0 to 11.0 eV, −2.0 to −1.0 eV, and 0.0 to 1.0 debye, respectively (Figure 1). According to the generation procedure, each model predicts that all of the generated structures will exhibit properties within the respective target ranges, whereas the histogram represents the distribution of true values as characterized at the

https://doi.org/10.1021/acs.jpca.1c08191
J. Phys. Chem. A 2022, 126, 333−340

xTB level. Each histogram thus represents a simultaneous test of the predictive ability of each model as well as its generative capability.

For VIP and EA, the chosen property ranges are not found within the training data at all, whereas for DM a poorly represented range was instead selected due to the lower bound on DM values and the long tail for compounds with very high DM in the training data. The sampling technique proves very effective for EA, with the mean and the majority of the sampling distribution (57%) falling within the target region. While the results for VIP and DM are not as favorable, both distributions undergo a clear shift toward the targeted region, with a high number of the generated structures (17 and 11%, respectively) fulfilling the target criterion in both situations. We also note for the case of DM that the lower bound on possible values may impact the number of generated structures in this regime. For all three properties, the model has learned enough chemical information from the initial training set alone to determine the relationship between the property of interest and the targeted structures.

Single-Property Active Learning. For property ranges that represent fundamentally different chemistries than those found in the training data, the model may not have learned the necessary structure−function relationships to effectively generate new structures with the targeted properties. In such a scenario, the proposed active learning strategy can be used to incrementally sample regions of the latent space to enrich the model's understanding of structure−function relationships. As a demonstration, we performed a generative search for structures exhibiting EA values between 1.0 and 2.0 eV, which is approximately one standard deviation higher than the mean EA of the training data but still in the interpolative regime (Figure 2). While 17% of structures sampled from the model are in the targeted EA range, this is only marginally higher than the training distribution and reflects limited specificity for high EA species. In contrast to the single-shot learning results for negative EA values presented in Figure 1B, it is also clear that the model has yet to resolve the structure−function relationship for predicting positive EA values even though it has some representation in the training data. Although the model has not yet learned a strong relationship between the chemical structure and the targeted EA range, the sampling still yields a large number of new structures exhibiting EA within or near the targeted range. Using the iterative approach outlined above, these new structures were incorporated into the training data to allow the network to resolve the functional relationships in the targeted EA region. After four iterations of sampling and retraining, the bulk of the sampled distribution shifted, with 30% of structures falling within the desired range. After nine iterations, the mean of the distribution shifted squarely within the 1−2 eV range and 35% of the sampled structures exhibited EA values within the target. Thus, even for single property optimization, the active learning approach is effective in generating samples that teach the model to understand and generate high EA compounds. Figure S5 illustrates the convergence of the generated property distribution to the target range as a function of iterations.

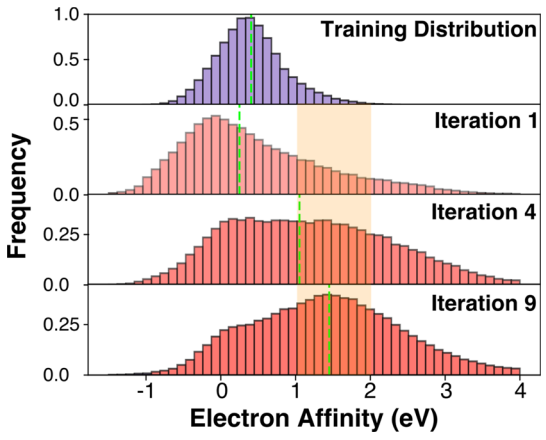

Figure 2. Property histograms for the model trained to predict EA. Training distribution is shown in purple (TOP), means are indicated by dashed green lines, and the target region of 1−2 eV is highlighted in yellow. Initially, the model has difficulty generating structures with the specified EA (iteration 1). After four iterations, the mean of the generated EAs has shifted to be within the target range. After nine iterations, the peak of the distribution is within the targeted range.

Multi-Target Optimization. While one property may be of primary interest in a particular molecular search (i.e., single-target optimization), there are often multiple properties that must be optimized simultaneously. This is often a much more difficult task, as these properties may have varying degrees of correlation and representation in the training data, an issue that is further compounded when considering the exponential growth of the property space with respect to the number of optimized properties. For instance, the challenge of multi-property optimization is apparent if we consider searching for compounds with EA between 1.5 and 4.0 eV, DM between 4.0 and 5.0 debye, and VIP above 10.0 eV. The DM and EA ranges are represented in the training data, and the property range for VIP may be individually sampled effectively, as the experiment in Figure 1 demonstrated. However, when considering all three properties in tandem, no training structures simultaneously exhibit this range of values. Moreover, attempting to sample structures from this region of the latent space is unsuccessful (Figure 3A), with none of the 100,000 sampled compounds falling within the desired property ranges. Figure 3A also demonstrates that simply retraining on new molecules optimized for individual properties would be ineffective, as their other properties are highly unlikely to fall within the desired range. However, after 8 iterations of retraining and resampling (Figure 3B), the sampled distribution shifted toward the multi-dimensional property target, with over 1600 target structures that simultaneously satisfy all three criteria being successfully generated. This demonstrates the potential for active learning as a framework to effectively fill in a model's chemical understanding, particularly in the case of multi-target extrapolative searching, where the increased dimensionality of the search space decreases the potential coverage of the training data. Figure S6 demonstrates the convergence of the overall generated property distribution toward the targeted ranges as a function of the number of iterations.

When performing any multi-objective optimization, there is the possibility of intrinsic trade-offs between the optimization targets. In the case of generative chemical models, these trade-offs are governed by the underlying structure−function relationships and will potentially vary dramatically across combinations of properties. The trade-offs associated with the current optimization of IP, EA, and DM are elucidated by the three pairs of Pareto front plots shown in Figures S7−S9 as a function of active learning iterations. After a single iteration, the targeted DM−VIP ranges and DM−EA ranges lie within

https://doi.org/10.1021/acs.jpca.1c08191
J. Phys. Chem. A 2022, 126, 333−340

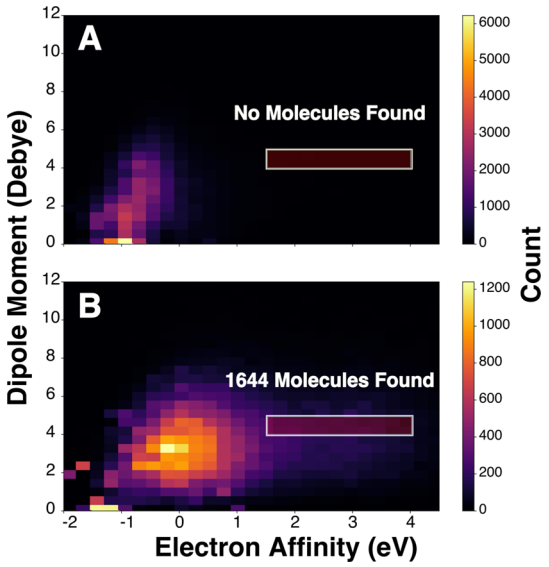

Figure 3. 2D property histogram for the model trained to predict VIP, EA, and DM and tasked with the targeted structure generation for these properties. For visualization, only compounds with VIP greater than 10.0 eV are considered. The targeted region, with DM between 4 and 5 debye and EA between 1.5 and 4.0 eV, is indicated with a box. The targeted range for VIP is extrapolated outside the training data, while the EA range has little representation and the DM range is very well represented. (A) Initially, the model is not effective in generating compounds that fulfil all three criteria together. (B) After eight iterations of the active learning procedure, the property distribution of the proposed structures has shifted to cover the targeted region and the model generates over 1600 structures fulfilling all three property criteria as validated at the xTB level.

the respective Pareto fronts, whereas the VIP−EA Pareto front has marginal penetration within the targeted range. Nevertheless, after three iterations all property combinations lie within their respective Pareto fronts. We conclude that the VIP−EA property combination exhibits the strongest trade-off for this particular optimization problem. Looking beyond the current study, we could envision pathological combinations of properties that might require more active learning iterations to discover satisfactory structures or even unphysical combinations of properties that would not be reachable with realistic molecules.

External Validation. In a practical scenario, the multi-target search discussed above would be the first step in a computational funnel to pare down the search space of viable molecules to a promising set for experimental study. However, we can further tighten this computational funnel through additional screening at the DFT level. Given the discrepancy between property calculations at the semi-empirical xTB and DFT levels, when selecting molecules from the final round of multi-target optimization for further screening we allowed for a soft-cutoff by adding ±20% of the target range to the property bounds in order to avoid screening out near-misses. To focus only on those compounds with the potential to be easily synthesized, we further reduced the list by screening out radicals, charged species, zwitterions, and structures with experimentally infeasible structures, such as those with linear oxygen chains of more than two atoms. This resulted in 307 candidate structures, which were then characterized at the ωB97X-D3/def2-TZVP level. After DFT characterizations, 22 structures passed the soft-cutoff criteria for all properties

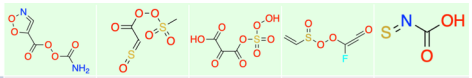

Vertical Ionization Potential (eV)	10.69	10.69	11.07	10.56	10.01
Electron Affinity (eV)	1.62	1.64	1.60	1.65	2.71
Dipole Moment (Debye)	4.80	4.60	4.18	4.24	4.11

Figure 4. Final five structures simultaneously achieving the desired VIP, EA, and DM values validated at the ωB97X-D3/def2-TZVP level. The model has learned to make heavy use of oxygen atoms to achieve a high ionization potential while simultaneously maintaining a high EA. The stability of these structures was confirmed by frequency calculations.

(Figure S2) and 5 passed the exact criteria for all properties (Figure 4). Inspecting the passing structures provides insight into the structure−function relationships that the model has learned to meet the targeted property ranges. We immediately note that all of the proposed structures are oxygen-rich. This is consistent with the high ionization potential target, which is promoted by including highly electronegative atoms and associated functional groups. Only one electronegative fluorine substitution is present in the final structures, although several examples occur in the structures that satisfy the relaxed criteria (Figure S2). We provisionally interpret the preference for oxygen substitutions over fluorine substitutions as being due to the interplay between the DM and ionization potential targets. In particular, we only observe isolated fluorine substitutions, which is consistent with the relatively high DM target. The high occurrence of oxygen, as well as nitrogen and fluorine, also promotes high EA, the second targeted property. Finally, we note that the model has also learned to avoid symmetric structures, which is necessary for the generation of molecules with a large DM.

As an additional demonstration, we conducted an analogous study on multi-property optimization in an interpolative regime with respect to the training data. The targeted property values (VIP ∈ [6.0, 7.0 eV], EA ∈ [0.5, 1.0 eV], and DM ∈ [4, 5 D]), had some representation in the training data (∼160 samples), but were shifted from the mean of each property value (Figure S3). VIP was shifted downward by approximately two standard deviations, while EA was shifted upward by approximately one standard deviation. DM was not shifted relative to the mean to ensure some representation of the selected property ranges in the training data. After 6 iterations of the active learning-based retraining, 1599 novel structures (i.e., structures that were not present in the training data) were sampled that satisfied the targeted property ranges at the xTB level, and of these 16 matched all property ranges after validation at the ωB97X-D3/def2-TZVP level (Figure S4). Intuitively, fewer cycles of retraining were required to find more matching structures because the property ranges already exhibited some representation with the training data set.

CONCLUSIONS

Multi-objective chemical optimization presents unique challenges compared with single-objective optimization, such as achieving simultaneous specificity for multiple targets, and data sparsity due to the increased dimensionality of the property search space. We have demonstrated that generative models can be coupled with active learning-based retraining to predict novel structures designed to have specific properties, even when these properties may not be observed in the training

https://doi.org/10.1021/acs.jpca.1c08191
J. Phys. Chem. A 2022, 126, 333−340

data. For difficult targets, particularly multi-objective targets, a generative chemical model can learn the prerequisite chemistries by iteratively retraining on compounds it proposes that are similar to those that are desired. In this way, the model generates its own nascent structure−function relationships that it refines by sampling the predicted structures. We have shown the ability of this method to propose compounds with specific VIP, EA, and DM individually and simultaneously and anticipate its utility for other sets of chemical properties.

Nevertheless, we envision several potential challenges in generalizing this approach to more urgent optimization problems. In particular, the properties used in this study were mainly selected out of convenience due to their ease of calculation and data abundance. For data scarce scenarios, or situations where querying an oracle involves performing an experiment or a more expensive simulation, generative models will likely remain prohibitively costly to train for multi-objective searches, even with active learning strategies. Additionally, the selected properties in the current study exhibited relatively modest trade-offs in our multi-objective optimization experiments. When adapting this approach to other combinations of properties or property ranges, it is possible that more severe trade-offs could be encountered. We also note that the quality of training data is a critical factor to ensure that the properties of the proposed molecules accurately match their true values. The discrepancy noted between the property values at the xTB and DFT levels could be relieved by using an auxiliary difference model that predicts the difference between the low- and high-accuracy computational methods and optimizing with respect to this variable instead. This may allow for more efficient sampling and fewer iterations of the active learning procedure. Exploiting higher order correlations between properties by employing a more complex predictor network along with a suitable algorithm for searching within the latent space also holds the potential for reducing training data requirements in data scarce scenarios.

■ ASSOCIATED CONTENT

Supporting Information

The Supporting Information is available free of charge at https://pubs.acs.org/doi/10.1021/acs.jpca.1c08191.

Model architecture schematic, additional training details, chemical structures from generative experiments, and Pareto fronts (PDF)

■ AUTHOR INFORMATION

Corresponding Author

Brett M. Savoie − *Charles D. Davidson School of Chemical Engineering, Purdue University, West Lafayette, Indiana 47906, United States;* orcid.org/0000-0002-7039-4039; Email: bsavoie@purdue.edu

Authors

Nicolae C. Iovanac − *Charles D. Davidson School of Chemical Engineering, Purdue University, West Lafayette, Indiana 47906, United States*

Robert MacKnight − *Charles D. Davidson School of Chemical Engineering, Purdue University, West Lafayette, Indiana 47906, United States*

Complete contact information is available at: https://pubs.acs.org/10.1021/acs.jpca.1c08191

Notes

The authors declare no competing financial interest.

Training data and xTB calculated properties for all sampled structures from all of the generative active learning have been uploaded to Zenodo with a persistent DOI (10.5281/zenodo.5512896). Sample model training scripts are available on GitHub (https://github.com/savoie-group/Actively-Searching).

■ ACKNOWLEDGMENTS

The work performed by N.C.I. and B.M.S. was made possible by the National Science Foundation (NSF) Division of Chemical, Bioengineering, Environmental, and Transport Systems (CBET) through support provided by the Electrochemical Systems Program (Grant number: 2045887-CBET, Program Manager: Dr. Carol Read). B.M.S. also acknowledges support from the Dreyfus Program for Machine Learning in the Chemical Sciences and Engineering during this project.

■ REFERENCES

(1) Rendall, R.; Castillo, I.; Lu, B.; Colegrove, B.; Broadway, M.; Chiang, L. H.; Reis, M. S. Image-Based Manufacturing Analytics: Improving the Accuracy of an Industrial Pellet Classification System Using Deep Neural Networks. *Chemom. Intell. Lab. Syst.* **2018**, *180*, 26−35.

(2) Esteva, A.; Kuprel, B.; Novoa, R. A.; Ko, J.; Swetter, S. M.; Blau, H. M.; Thrun, S. Dermatologist-Level Classification of Skin Cancer with Deep Neural Networks. *Nature* **2017**, *542*, 115−118.

(3) Hansen, K.; Biegler, F.; Ramakrishnan, R.; Pronobis, W.; von Lilienfeld, O. A.; Müller, K.-R.; Tkatchenko, A. Machine Learning Predictions of Molecular Properties: Accurate Many-Body Potentials and Nonlocality in Chemical Space. *J. Phys. Chem. Lett.* **2015**, *6*, 2326−2331.

(4) Faber, F. A.; Hutchison, L.; Huang, B.; Gilmer, J.; Schoenholz, S. S.; Dahl, G. E.; Vinyals, O.; Kearnes, S.; Riley, P. F.; von Lilienfeld, O. A. Prediction Errors of Molecular Machine Learning Models Lower than Hybrid DFT Error. *J. Chem. Theory Comput.* **2017**, *13*, 5255−5264.

(5) Altarazi, S.; Allaf, R.; Alhindawi, F. Machine Learning Models for Predicting and Classifying the Tensile Strength of Polymeric Films Fabricated via Different Production Processes. *Materials* **2019**, *12*, 1475.

(6) Matos, M. A. S.; Pinho, S. T.; Tagarielli, V. L. Application of Machine Learning to Predict the Multiaxial Strain-Sensing Response of CNT-Polymer Composites. *Carbon* **2019**, *146*, 265−275.

(7) Xie, T.; Grossman, J. C. Crystal Graph Convolutional Neural Networks for an Accurate and Interpretable Prediction of Material Properties. *Phys. Rev. Lett.* **2018**, *120*, 145301.

(8) Sanchez-Lengeling, B.; Outeiral, C.; Guimaraes, G. L.; Aspuru-Guzik, A. Optimizing Distributions over Molecular Space. An Objective-Reinforced Generative Adversarial Network for Inverse-Design Chemistry (ORGANIC), **2017**. ChemRxiv:10.26434/chemrxiv.5309668.v3.

(9) De Cao, N.; Kipf, T. MolGAN: An Implicit Generative Model for Small Molecular Graphs, **2018**. arXiv:1805.11973.

(10) Kadurin, A.; Nikolenko, S.; Khrabrov, K.; Aliper, A.; Zhavoronkov, A. DruGAN: An Advanced Generative Adversarial Autoencoder Model for de Novo Generation of New Molecules with Desired Molecular Properties in Silico. *Mol. Pharm.* **2017**, *14*, 3098−3104.

(11) Maziarka, Ł.; Pocha, A.; Kaczmarczyk, J.; Rataj, K.; Danel, T.; Warchoł, M. Mol-CycleGAN: A Generative Model for Molecular Optimization. *J. Cheminf.* **2020**, *12*, 2.

(12) Sattarov, B.; Baskin, I. I.; Horvath, D.; Marcou, G.; Bjerrum, E. J.; Varnek, A. De Novo Molecular Design by Combining Deep Autoencoder Recurrent Neural Networks with Generative Topographic Mapping. *J. Chem. Inf. Model.* **2019**, *59*, 1182.

(13) Polykovskiy, D.; Zhebrak, A.; Vetrov, D.; Ivanenkov, Y.; Aladinskiy, V.; Mamoshina, P.; Bozdaganyan, M.; Aliper, A.; Zhavoronkov, A.; Kadurin, A. Entangled Conditional Adversarial Autoencoder for de Novo Drug Discovery. *Mol. Pharm.* **2018**, *15*, 4398−4405.
(14) Lim, J.; Ryu, S.; Kim, J. W.; Kim, W. Y. Molecular Generative Model Based on Conditional Variational Autoencoder for de Novo Molecular Design. *J. Cheminf.* **2018**, *10*, 31.
(15) Kadurin, A.; Aliper, A.; Kazennov, A.; Mamoshina, P.; Vanhaelen, Q.; Khrabrov, K.; Zhavoronkov, A. The Cornucopia of Meaningful Leads: Applying Deep Adversarial Autoencoders for New Molecule Development in Oncology. *Oncotarget* **2016**, *8*, 10883−10890.
(16) Blaschke, T.; Olivecrona, M.; Engkvist, O.; Bajorath, J.; Chen, H. Application of Generative Autoencoder in De Novo Molecular Design. *Mol. Inf.* **2018**, *37*, 1700123.
(17) Griffiths, R.-R.; Hernández-Lobato, J. M. Constrained Bayesian Optimization for Automatic Chemical Design Using Variational Autoencoders. *Chem. Sci.* **2020**, *11*, 577−586.
(18) Hong, S. H.; Ryu, S.; Lim, J.; Kim, W. Y. Molecular Generative Model Based on an Adversarially Regularized Autoencoder. *J. Chem. Inf. Model.* **2020**, *60*, 29−36.
(19) 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.
(20) Jin, W.; Barzilay, R.; Jaakkola, T. Junction Tree Variational Autoencoder for Molecular Graph Generation, **2018**. arXiv:1802.04364. ArXiv [cs.LG].
(21) Li, Y.; Vinyals, O.; Dyer, C.; Pascanu, R.; Battaglia, P. Learning Deep Generative Models of Graphs, **2018**. arXiv:1803.03324.
(22) Simonovsky, M.; Komodakis, N. GraphVAE: Towards Generation of Small Graphs Using Variational Autoencoders, **2018**. ArXiv [cs.LG] arXiv:1802.03480.
(23) 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*, 241735.
(24) Wu, S.; Kondo, Y.; Kakimoto, M.-a.; Yang, B.; Yamada, H.; Kuwajima, I.; Lambard, G.; Hongo, K.; Xu, Y.; Shiomi, J.; Schick, C.; Morikawa, J.; Yoshida, R. Machine-Learning-Assisted Discovery of Polymers with High Thermal Conductivity Using a Molecular Design Algorithm. *npj Comput. Mater.* **2019**, *5*, 1−11.
(25) Janet, J. P.; Ramesh, S.; Duan, C.; Kulik, H. J. Accurate Multiobjective Design in a Space of Millions of Transition Metal Complexes with Neural-Network-Driven Efficient Global Optimization. *ACS Cent. Sci.* **2020**, *6*, 513−524.
(26) Domenico, A.; Nicola, G.; Daniela, T.; Fulvio, C.; Nicola, A.; Orazio, N. De Novo Drug Design of Targeted Chemical Libraries Based on Artificial Intelligence and Pair-Based Multiobjective Optimization. *J. Chem. Inf. Model.* **2020**, *60*, 4582.
(27) Ståhl, N.; Falkman, G.; Karlsson, A.; Mathiason, G.; Boström, J. Deep Reinforcement Learning for Multiparameter Optimization in de Novo Drug Design. *J. Chem. Inf. Model.* **2019**, *59*, 3166−3176.
(28) Nigam, A.; Pollice, R.; Krenn, M.; Gomes, G. d. P.; Aspuru-Guzik, A. Beyond Generative Models: Superfast Traversal, Optimization, Novelty, Exploration and Discovery (STONED) Algorithm for Molecules Using SELFIES. *Chem. Sci.* **2021**, *12*, 7079.
(29) Zhou, Z.; Kearnes, S.; Li, L.; Zare, R. N.; Riley, P. Optimization of Molecules via Deep Reinforcement Learning. *Sci. Rep.* **2019**, *9*, 10752.
(30) Settles, B. Active Learning Literature Survey. *Technical Report*; University of Wisconsin-Madison Department of Computer Sciences, 2009.
(31) Konze, K. D.; Bos, P. H.; Dahlgren, M. K.; Leswing, K.; Tubert-Brohman, I.; Bortolato, A.; Robbason, B.; Abel, R.; Bhat, S. Reaction-Based Enumeration, Active Learning, and Free Energy Calculations To Rapidly Explore Synthetically Tractable Chemical Space and Optimize Potency of Cyclin-Dependent Kinase 2 Inhibitors. *J. Chem. Inf. Model.* **2019**, *59*, 3782−3793.
(32) Ghanakota, P.; Bos, P. H.; Konze, K. D.; Staker, J.; Marques, G.; Marshall, K.; Leswing, K.; Abel, R.; Bhat, S. Combining Cloud-Based Free-Energy Calculations, Synthetically Aware Enumerations, and Goal-Directed Generative Machine Learning for Rapid Large-Scale Chemical Exploration and Optimization. *J. Chem. Inf. Model.* **2020**, *60*, 4311−4325.
(33) Kim, C.; Chandrasekaran, A.; Jha, A.; Ramprasad, R. Active-Learning and Materials Design: The Example of High Glass Transition Temperature Polymers. *MRS Commun.* **2019**, *9*, 860−866.
(34) Zhong, M.; Tran, K.; Min, Y.; Wang, C.; Wang, Z.; Dinh, C.-T.; De Luna, P.; Yu, Z.; Rasouli, A. S.; Brodersen, P.; Sun, S.; Voznyy, O.; Tan, C.-S.; Askerka, M.; Che, F.; Liu, M.; Seifitokaldani, A.; Pang, Y.; Lo, S.-C.; Ip, A.; Ulissi, Z.; Sargent, E. H. Accelerated Discovery of CO 2 Electrocatalysts Using Active Machine Learning. *Nature* **2020**, *581*, 178−183.
(35) Sterling, T.; Irwin, J. J. ZINC 15 − Ligand Discovery for Everyone. *J. Chem. Inf. Model.* **2015**, *55*, 2324−2337.
(36) Bannwarth, C.; Ehlert, S.; Grimme, S. GFN2-XTB—An Accurate and Broadly Parametrized Self-Consistent Tight-Binding Quantum Chemical Method with Multipole Electrostatics and Density-Dependent Dispersion Contributions. *J. Chem. Theory Comput.* **2019**, *15*, 1652−1671.
(37) Adamo, C.; Barone, V. Toward Reliable Density Functional Methods without Adjustable Parameters: The PBE0 Model. *J. Chem. Phys.* **1999**, *110*, 6158−6170.
(38) Neese, F. Software Update: The ORCA Program System, Version 4.0. *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **2018**, *8*, No. e1327.
(39) Kusner, M. J.; Paige, B.; Hernández-Lobato, J. M. Grammar Variational Autoencoder. *ICML'17: Proceedings of the 34th International Conference on Machine Learning*, 2017; Vol. 70, pp 1945−1954.
(40) Iovanac, N. C.; Savoie, B. M. Simpler Is Better: How Linear Prediction Tasks Improve Transfer Learning in Chemical Autoencoders. *J. Phys. Chem. A* **2020**, *124*, 3679−3685.
(41) Chollet, F. K. *Keras*, 2015. https://github.com/fchollet/keras.
(42) Abadi, M.; et al. TensorFlow: A System for Large-Scale Machine Learning. *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, 2016; pp 265−283.
(43) Brochu, E.; Cora, V. M.; de Freitas, N. A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning, **2010**. arXiv:1012.2599. ArXiv [cs.LG].
(44) Shmilovich, K.; Mansbach, R. A.; Sidky, H.; Dunne, O. E.; Panda, S. S.; Tovar, J. D.; Ferguson, A. L. Discovery of Self-Assembling π-Conjugated Peptides by Active Learning-Directed Coarse-Grained Molecular Simulation. *J. Phys. Chem. B* **2020**, *124*, 3873−3891.
(45) Ramakrishnan, R.; Dral, P. O.; Rupp, M.; von Lilienfeld, O. A. Quantum Chemistry Structures and Properties of 134 Kilo Molecules. *Sci. Data* **2014**, *1*, 140022.
(46) Iovanac, N. C.; Savoie, B. M. Improving the Generative Performance of Chemical Autoencoders through Transfer Learning. *Mach. Learn.: Sci. Technol.* **2020**, *1*, 045010.

https://doi.org/10.1021/acs.jpca.1c08191
J. Phys. Chem. A 2022, 126, 333−340