

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

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

Charles D. Davidson School of Chemical Engineering, 480 Stadium Mall Drive, Purdue University, West Lafayette, IN 47906, USA.

Abstract: Combining quantum chemistry characterizations with generative machine learning models has the potential to accelerate molecular searches in chemical space. 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 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 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 to obtain the true values, and these 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 unknown areas of chemical space and shifts the distribution of generated compounds towards 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 ω B97X-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 analytics¹ and cancer detection.² In recent years, it has proved particularly successful in the chemical sciences, where ML has been used to predict interatomic potentials³, quantum chemical properties⁴, 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 chemical optimization methods by providing reasonable chemical suggestions for more expensive experimental synthesis and characterization. Generative adversarial networks^{8,9,10,11} (GANs) and various formulations of autoencoder networks^{12,13,14,15,16,17,18} have emerged as some of the more popular frameworks for generative machine-learning based chemical design. These methods often provide for a predictive aspect which allows suggestions to be biased towards compounds with particular properties.¹⁹ Much effort has been directed to solving issues related to the ability of these models to generate valid chemistries^{20,9,21,22}, and they have been successfully demonstrated in generating compounds with specific properties such as bandgap²³ 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 challenging

because they will typically be 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 solubility and redox potential in the design of transition metal complexes for redox flow batteries using efficient global optimization to explore an enumerated space of 2.8 million candidate complexes.²⁵ Domenico *et al.* utilize 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, are 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 optimized molecular weight, logP, and polar surface area.²⁷ Nigram *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 logP and 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. LogP, 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 datasets 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 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 targeted regions of chemical space. Generative chemical models have the unique feature that syntactically valid outputs are guaranteed to belong within chemical space, meaning that they are suitable samples for further model training. By sampling underrepresented 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 the 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 We propose formulating the entire goal of multi-target chemical optimization as an active learning problem. Rather than attempting to determine optimal regions of chemical space to sample or train on, we query to model to obtain its suggestions for compounds with the desired properties. These compounds are then screened via semi-empirical 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 the opportunity to continuously improve a model's ability to target compounds with underrepresented 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 database³³, we develop our own dataset of quantum chemical properties including vertical ionization potential (VIP), electron affinity (EA), and dipole moment (DM) calculated at the semi-empirical level for training. 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 ω B97X-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

Datasets. All models are initially trained and evaluated using structures from the ZINC15 dataset. This dataset contains 3D structural data for hundreds of millions of small molecules, from which we have chosen a subset of 250,000 compounds with molecular weight between 200 and 500 Daltons and logP between -1 and 5. These compounds were subjected to geometry optimization and electronic structure calculation with GFN2-xTB³⁴ (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. 80 percent were utilized for training, with the remaining 20 percent withheld for validation. Additional structures generated during the active learning step were subjected to the same property calculation methods to expand the dataset. In order to validate

the properties of the structures generated in the multi-objective active learning study, xTB optimized geometries are used as input to a geometry optimization at the ω B97X-D3/def2-TZVP level to determine the dipole moment of the neutral species.³⁵ Using the optimized neutral geometries, the DFT calculations are repeated for the cation and anion to calculate the vertical ionization potential and electron affinity, respectively. All DFT calculations were conducted using Orca 4.0.1.³⁶

Machine Learning Architecture. Three models were developed for single target predictions and one model was developed for multi-property prediction tasks. We utilize the grammar variational autoencoder³⁷ (GVAE) to achieve generation of molecules with targeted properties, with the alteration of using a single linear predictor layer so that properties tend to vary linearly along the latent dimensions.³⁸ 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. KL divergence loss was scaled by 750, and the categorical cross-entropy loss associated with encoding and decoding was decayed from 50 to 1 during training according to a sigmoid function. All properties were normalized to fall within a range of -20 to 20. The normalization and scaling factors were selected to balance property prediction accuracy, encoding and decoding accuracy, and the ability to decode novel structures from arbitrary latent points. Additional training details are provided in the SI.

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 those properties will vary linearly

along the principal components of the latent encodings.³⁸ Compounds with specific properties can then be generated by targeting regions of the latent space based on univariate linear regression between the property of interest and the position along one of the principle components. In the case of multi-property models, correlation between the properties may lead to latent space organization not being exactly orthogonal. To find the direction to sample, we linearly regress the angles that the points must be rotated by to maximize the R^2 between position along a particular principal component and the property of interest. This allows us to continue to utilize a simple linear regression to target specific properties. All other dimensions are sampled normally with mean and standard deviation determined by the training data.

Active Learning Technique. Depending on the nature of the property and the range for which it is sampled, and particularly for extrapolative property searches, the model may not be able to generate structures with properties that match those suggested by the regression outlined above. To circumvent this, structures are sampled from the regions of latent space that the model predicts to be of interest, which are then used for retraining the model. In each iteration, 100,000 unique structures are sampled from the model and the canonical SMILES of these structures are checked against the training and validation datasets to ensure novelty of the generated structures. Those novel structures are then subjected to xTB calculations for characterization. As the sampling routine is not perfect, these compounds are filtered to ensure they fall within the desired range before usage in retraining. In situations where no compounds fall within the desired range, we instead select molecules with properties that fall 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 are closer to the targeted range than the original training data. It was

found that simply retraining the model 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 are randomly sampled and added to this new dataset until reaching 100,000 total training structures. This dataset size was found to train effectively using the same hyperparameters initially used in training. This routine was repeated until the desired number of structures with the targeted property ranges was obtained.

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 dataset³⁹, namely internal energy, zero-point vibrational energy, and HOMO-LUMO gap.⁴⁰ To examine the generality of this approach, 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




Figure 1: Property histograms for molecules generated by models trained on (A) vertical ionization potential, (B) electron affinity, and (C) dipole moment. 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 extrapolating to compounds with property values not observed in the training data, shown in orange.

eV, and 0.0 to 1.0 Debye, respectively (**Fig. 1**). 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 extreme, both distributions undergo a clear shift towards the targeted region, with a high number of 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. 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 (**Fig. 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. Although the model has not yet learned a strong relationship between 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 in the methods section, these new structures were incorporated into the training data to allow the network to resolve the functional relationships in the targeted EA region.

After 4 iterations of sampling and retraining, the bulk of the sampled distribution shifted, with 30% of structures falling within the desired range. After 9 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 teaching the model the missing chemistries it needs to understand and generate high EA compounds.

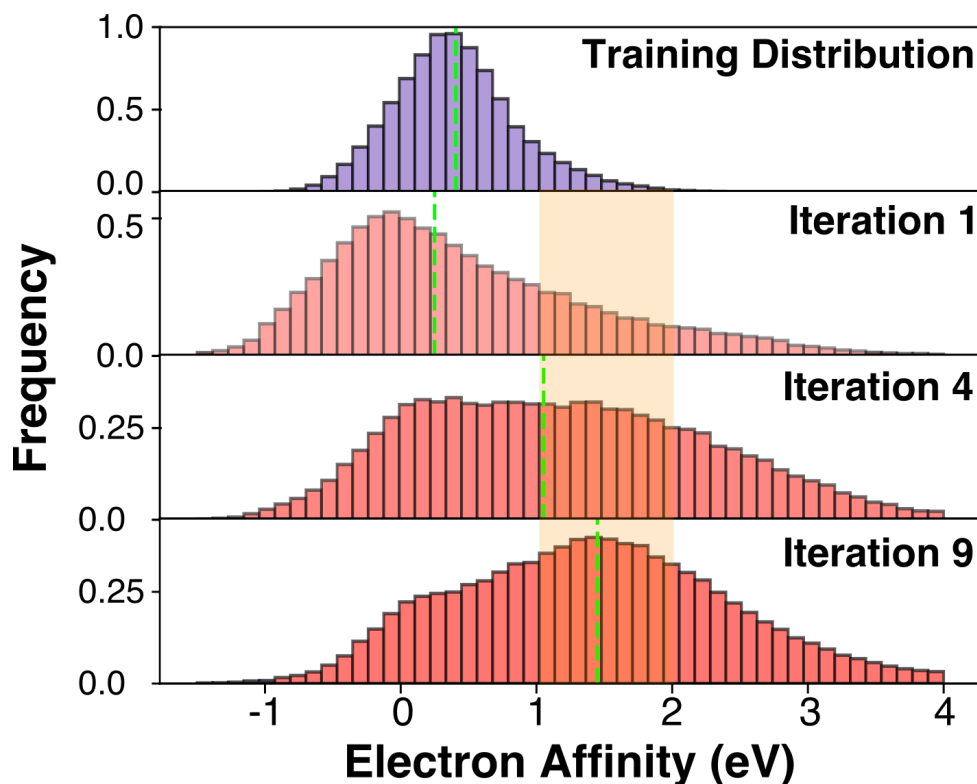


Figure 2: Property histograms for model trained to predict electron affinity. Training distribution is shown in purple (TOP), means are indicated by dashed green line, and the target region of 1-2 eV is highlighted in yellow. Initially, the model has difficulty generating structures with the specified EA (Iteration 1). After 4 and finally 9 iterations, the mean of the generated EAs has shifted to be within the target range, where the distribution also peaks.

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 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 1.5-4.0 eV, DM between 4.0-

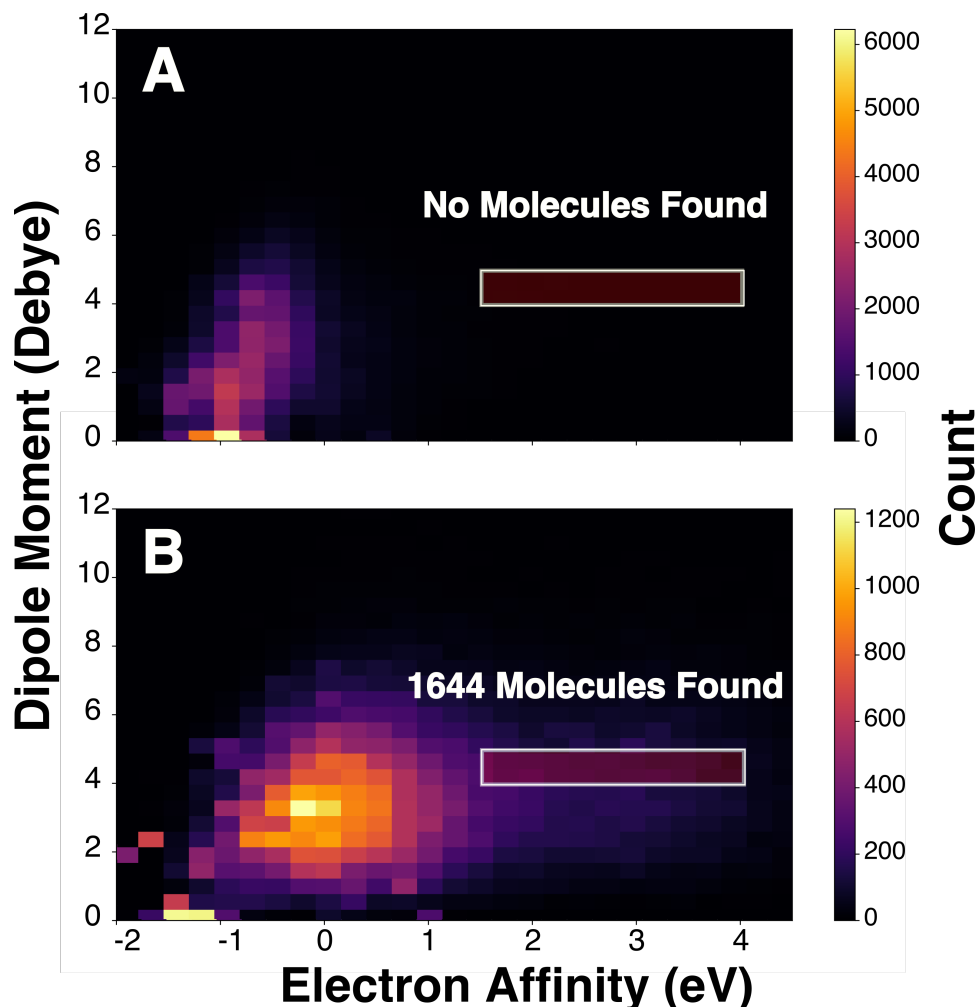


Figure 3: 2D property histogram for model trained to predict VIP, EA, and DM, and tasked with targeted structure generation for these properties. For visualization, only compounds with VIP greater than 10.0 eV are considered. The targeted region, with DM between 4-5 Debye and EA between 1.5-4.0 eV, is indicated with a box. Ionization potential is extrapolated beyond the training data, while the electron affinity range has little representation and the dipole moment range is very well represented. Initially, (A) the model is not effective in generating compounds that fulfil all three criteria together. After 8 iterations of the active learning procedure (B), the property distribution of proposed structures has shifted to cover the targeted region and the model is now capable of proposing over 1600 structures fulfilling all three property criteria.

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 (**Fig. 3A**), with none of 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 (**Fig. 3B**), the sampling distribution shifted towards the multi-dimensional property target, with over 1600 target structures 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 potential coverage of the training data.

External Validation. In a practical scenario, the active learning-based search procedure 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 a higher level of theory. Given the discrepancy between property calculations at the semi-empirical xTB and DFT levels, when selecting molecules 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 ω B97X-D3/def2-TZVP level. After

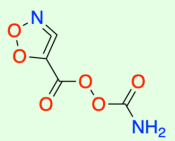
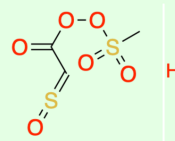
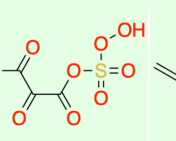
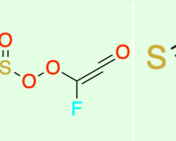
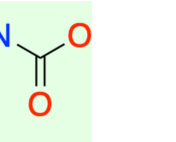
					
Vertical Ionization Potential (eV)	10.69	10.69	11.07	10.56	10.01
Electron Affinity (eV)	1.62	1.64	1.60	1.65	2.71
Dipole Moment (Debye)	4.80	4.60	4.18	4.24	4.11

Figure 4: Five structures simultaneously achieving targeted ionization potential, electron affinity, and dipole moment values validated at the ω B97X-D3/def2-TZVP level. The model has learned to make use of oxygen atoms to achieve a high ionization potential while simultaneously maintaining a high electron affinity and breaking symmetry to achieve the dipole moment target.

DFT characterizations, 22 structures passed the soft-cutoff criteria for all properties (**Fig. S1**), and 5 passed the exact criteria for all properties (**Fig. 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 (**Fig. S1**). We provisionally interpret the preference for oxygen substitutions over fluorine substitutions as being due to the interplay between the dipole moment and ionization potential targets. In particular, we only observe isolated fluorine substitutions, which is consistent with the relatively high dipole moment target. The high occurrence of oxygen, as well as nitrogen and fluorine, also promotes high electron affinity, 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 dipole moment.

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 ($VIP \in [6.0, 7.0 \text{eV}]$, $EA \in [0.5, 1.0 \text{eV}]$, and $DM \in [4, 5 \text{D}]$), had some representation in the training data (~ 160 samples), but were shifted from the mean of each property value (**Fig. S2**). Vertical ionization potential was shifted downwards by approximately two standard deviations, while electron affinity was shifted upwards by approximately one standard deviation. Dipole moment 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 were sampled that satisfy the targeted property ranges at the xTB level, and of these 16 matched all property ranges after validation at the ω B97X-D3/def2-TZVP level (**Fig. S3**). Intuitively, fewer cycles of retraining were required to find more matching structures, since the property ranges already exhibited some representation with the training dataset.

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 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 predicted structures. We have shown the ability of this method to propose compounds with specific vertical ionization potential, electron affinity, and dipole

moment individually and simultaneously, and anticipate its utility for other sets of chemical properties. We also note that the quality of training data is a critical factor to ensure the properties of proposed molecules accurately match their true values. The discrepancy noted between 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.

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. <https://doi.org/10.1016/j.chemolab.2018.07.001>.
- (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* (7639), 115–118. <https://doi.org/10.1038/nature21056>.
- (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* (12), 2326–2331. <https://doi.org/10.1021/acs.jpcllett.5b00831>.
- (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* (11), 5255–5264. <https://doi.org/10.1021/acs.jctc.7b00577>.
- (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* (9), 1475. <https://doi.org/10.3390/ma12091475>.
- (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. <https://doi.org/10.1016/j.carbon.2019.02.001>.
- (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* (14), 145301. <https://doi.org/10.1103/PhysRevLett.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). *figshare* **2017**. <https://doi.org/10.26434/chemrxiv.5309668.v3>.
- (9) De Cao, N.; Kipf, T. MolGAN: An Implicit Generative Model for Small Molecular Graphs. *ArXiv180511973 Cs Stat* **2018**.
- (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* (9), 3098–3104. <https://doi.org/10.1021/acs.molpharmaceut.7b00346>.
- (11) Łukasz, M.; Link to external site, this link will open in a new window; Agnieszka, P.; Link to external site, this link will open in a new window; Kaczmarczyk, J.; Link to external site, this link will open in a new window; Krzysztof, R.; Link to external site, this link will open in a new window; Tomasz, D.; Link to external site, this link will open in a new window; Michał, W.; Link to external site, this link will open in a new window. Mol-CycleGAN: A Generative Model for Molecular Optimization. *J. Cheminformatics Lond.* **2020**, *12* (1). <http://dx.doi.org/10.1186/s13321-019-0404-1>.
- (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**.
<https://doi.org/10.1021/acs.jcim.8b00751>.
- (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* (10), 4398–4405. <https://doi.org/10.1021/acs.molpharmaceut.8b00839>.
- (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. Cheminformatics* **2018**, *10* (1), 31. <https://doi.org/10.1186/s13321-018-0286-7>.
- (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* (7), 10883–10890. <https://doi.org/10.18632/oncotarget.14073>.
- (16) Blaschke, T.; Olivecrona, M.; Engkvist, O.; Bajorath, J.; Chen, H. Application of Generative Autoencoder in De Novo Molecular Design. *Mol. Inform.* **2018**, *37* (1–2), 1700123. <https://doi.org/10.1002/minf.201700123>.
- (17) Griffiths, R.-R.; Miguel Hernández-Lobato, J. Constrained Bayesian Optimization for Automatic Chemical Design Using Variational Autoencoders. *Chem. Sci.* **2020**, *11* (2), 577–586. <https://doi.org/10.1039/C9SC04026A>.
- (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* (1), 29–36. <https://doi.org/10.1021/acs.jcim.9b00694>.
- (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* (2), 268–276. <https://doi.org/10.1021/acscentsci.7b00572>.
- (20) Jin, W.; Barzilay, R.; Jaakkola, T. Junction Tree Variational Autoencoder for Molecular Graph Generation. *ArXiv180204364 Cs Stat* **2018**.
- (21) Li, Y.; Vinyals, O.; Dyer, C.; Pascanu, R.; Battaglia, P. Learning Deep Generative Models of Graphs. *ArXiv180303324 Cs Stat* **2018**.
- (22) Simonovsky, M.; Komodakis, N. GraphVAE: Towards Generation of Small Graphs Using Variational Autoencoders. *ArXiv180203480 Cs* **2018**.
- (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* (24), 241735. <https://doi.org/10.1063/1.5023563>.
- (24) Wu, S.; Kondo, Y.; Kakimoto, M.; 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), 1–11. <https://doi.org/10.1038/s41524-019-0203-2>.
- (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* (4), 513–524. <https://doi.org/10.1021/acscentsci.0c00026>.

- (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**. <https://doi.org/10.1021/acs.jcim.0c00517>.
- (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* (7), 3166–3176. <https://doi.org/10.1021/acs.jcim.9b00325>.
- (28) Nigam, A.; Pollice, R.; Krenn, M.; Passos Gomes, G. dos; Aspuru-Guzik, A. Beyond Generative Models: Superfast Traversal, Optimization, Novelty, Exploration and Discovery (STONED) Algorithm for Molecules Using SELFIES. *Chem. Sci.* **2021**. <https://doi.org/10.1039/D1SC00231G>.
- (29) Zhou, Z.; Kearnes, S.; Li, L.; Zare, R. N.; Riley, P. Optimization of Molecules via Deep Reinforcement Learning. *Sci. Rep.* **2019**, *9* (1), 1–10. <https://doi.org/10.1038/s41598-019-47148-x>.
- (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* (9), 3782–3793. <https://doi.org/10.1021/acs.jcim.9b00367>.
- (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* (9), 4311–4325. <https://doi.org/10.1021/acs.jcim.0c00120>.
- (33) Sterling, T.; Irwin, J. J. ZINC 15 – Ligand Discovery for Everyone. *J. Chem. Inf. Model.* **2015**, *55* (11), 2324–2337. <https://doi.org/10.1021/acs.jcim.5b00559>.
- (34) 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* (3), 1652–1671. <https://doi.org/10.1021/acs.jctc.8b01176>.
- (35) Adamo, C.; Barone, V. Toward Reliable Density Functional Methods without Adjustable Parameters: The PBE0 Model. *J. Chem. Phys.* **1999**, *110* (13), 6158–6170. <https://doi.org/10.1063/1.478522>.
- (36) Neese, F. Software Update: The ORCA Program System, Version 4.0. *Wiley Interdiscip. Rev. Comput. Mol. Sci.* **2018**, *8* (1), e1327. <https://doi.org/10.1002/wcms.1327>.
- (37) Kusner, M. J.; Paige, B.; Hernández-Lobato, J. M. Grammar Variational Autoencoder. *arXiv:1703.01925* **2017**.
- (38) 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* (18), 3679–3685. <https://doi.org/10.1021/acs.jpca.0c00042>.
- (39) 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. <https://doi.org/10.1038/sdata.2014.22>.

- (40) Iovanac, N. C.; Savoie, B. M. Improving the Generative Performance of Chemical Autoencoders through Transfer Learning. *Mach. Learn. Sci. Technol.* **2020**, *1* (4), 045010. <https://doi.org/10.1088/2632-2153/abae75>.