

PAPER • OPEN ACCESS

Improving the generative performance of chemical autoencoders through transfer learning

To cite this article: Nicolae C Iovanac and Brett M Savoie 2020 *Mach. Learn.: Sci. Technol.* 1 045010

View the [article online](#) for updates and enhancements.



ERRATUM

Erratum: Improving the generative performance of chemical autoencoders through transfer learning (2020 *Mach. Learn.: Sci. Technol.* **1** 045010)

OPEN ACCESS

RECEIVED
14 January 2021ACCEPTED FOR PUBLICATION
29 January 2021PUBLISHED
13 May 2021

Nicolae C Iovanac and Brett M Savoie

Charles D Davidson School of Chemical Engineering, Purdue University, 480 Stadium Mall Drive, West Lafayette, IN 47906, United States of America

Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



The following error was introduced during the production process. Figures 3 and 4 were printed incorrectly; the correct figures are reproduced here.


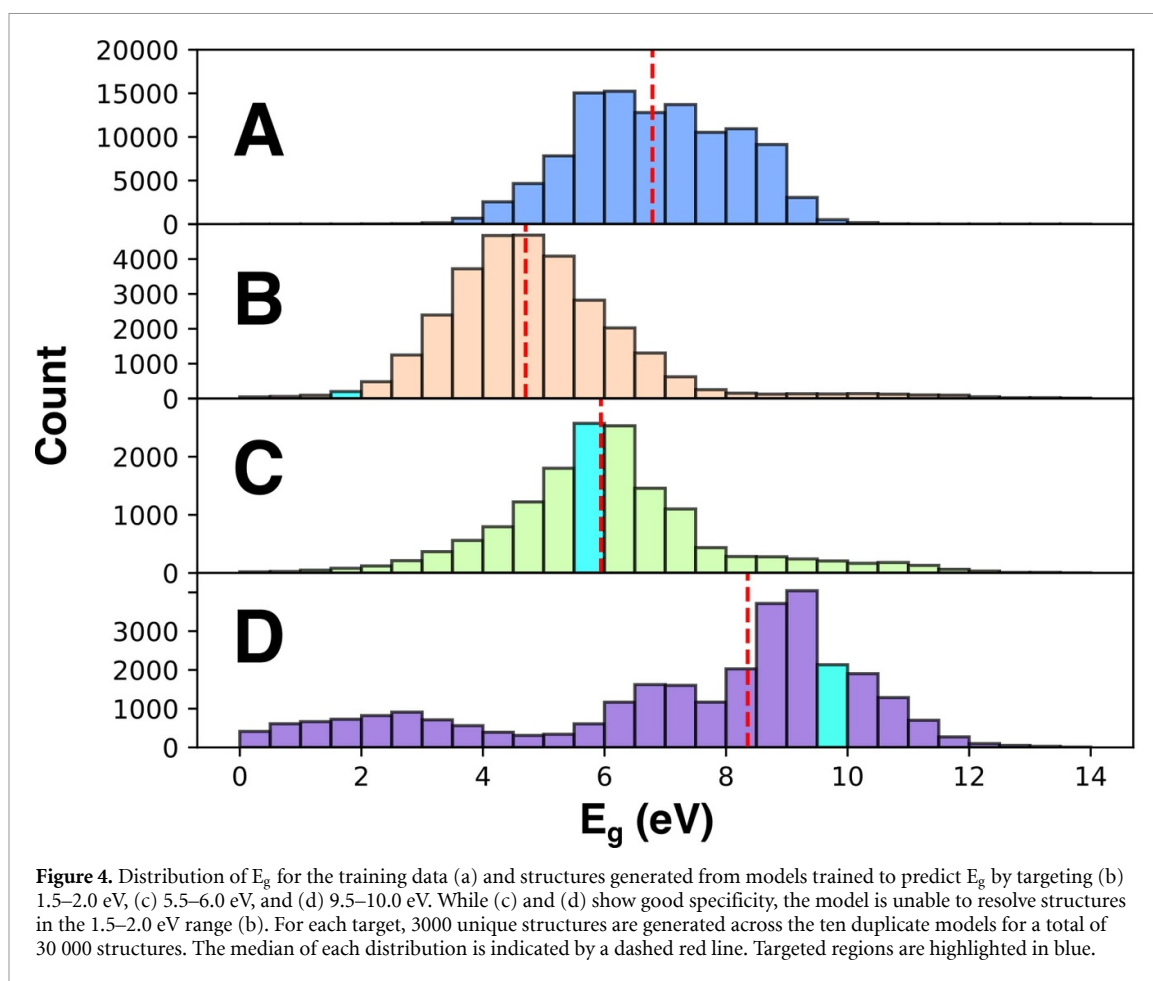


Figure 3. Distribution of U_0 for structures for the training set (a) compared with structures generated from extrapolation along the first principal component of models trained on (b) encoding and decoding alone, (c) an ancillary U_0 prediction task, (d) U_0 and ZPVE prediction tasks, and (e) U_0 , ZPVE, and E_g prediction tasks. For each paradigm, 3000 unique structures are generated across the ten duplicate models for a total of 30 000 structures. The mean of each distribution is denoted with a dashed red line and the largest $|U_0|$ value in the training data is indicated by the dashed green line. The percentage of generated compounds with $|U_0|$ greater than observed in the training data is shown on the right.



ORCID iD

Brett M Savoie  <https://orcid.org/0000-0002-7039-4039>



PAPER

OPEN ACCESS

RECEIVED
21 May 2020REVISED
24 July 2020ACCEPTED FOR PUBLICATION
12 August 2020PUBLISHED
8 October 2020

Improving the generative performance of chemical autoencoders through transfer learning

Nicolae C Iovanac and Brett M Savoie

Charles D Davidson School of Chemical Engineering, Purdue University, 480 Stadium Mall Drive, West Lafayette, IN 47906, United States of America

E-mail: bsavoie@purdue.edu**Keywords:** generative machine learning, transfer learning, chemical discovery

Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



Abstract

Generative models are a sub-class of machine learning models that are capable of generating new samples with a target set of properties. In chemical and materials applications, these new samples might be drug targets, novel semiconductors, or catalysts constrained to exhibit an application-specific set of properties. Given their potential to yield high-value targets from otherwise intractable design spaces, generative models are currently under intense study with respect to how predictions can be improved through changes in model architecture and data representation. Here we explore the potential of multi-task transfer learning as a complementary approach to improving the validity and property specificity of molecules generated by such models. We have compared baseline generative models trained on a single property prediction task against models trained on additional ancillary prediction tasks and observe a generic positive impact on the validity and specificity of the multi-task models. In particular, we observe that the validity of generated structures is strongly affected by whether or not the models have chemical property data, as opposed to only syntactic structural data, supplied during learning. We demonstrate this effect in both interpolative and extrapolative scenarios (i.e., where the generative targets are poorly represented in training data) for models trained to generate high energy structures and models trained to generate structures with targeted bandgaps within certain ranges. In both instances, the inclusion of additional chemical property data improves the ability of models to generate valid, unique structures with increased property specificity. This approach requires only minor alterations to existing generative models, in many cases leveraging prediction frameworks already native to these models. Additionally, the transfer learning strategy is complementary to ongoing efforts to improve model architectures and data representation and can foreseeably be stacked on top of these developments.

1. Introduction

The proliferation of machine learning (ML) research in the context of the chemical sciences has yielded powerful modeling paradigms for increasing the accuracy and reducing the cost of predicting the properties of molecules and materials. As this work reaches maturity, the so-called ‘forward-problem’ of predicting function from a given chemical structure is becoming more routine and with new datasets coming online more properties will soon become accessible to prediction using ML models. However, the ‘inverse-problem’ of finding an optimal set of structures under functional constraints is more directly relevant to chemical design and remains unsolved. Deep generative models are one class of ML modeling that is potentially capable of addressing the inverse problem by yielding exemplary structures from the same population as the training distribution. Such models have been effective in generating images that accurately match a desired caption [1], novel musical scores [2], and have recently been under intense study in chemical applications to discover new functional molecules and materials [3–8]. Common generative methods include variational

autoencoders [9–11], recurrent neural networks [12–14], generative adversarial networks [15–17], and adversarial autoencoders [18–20], among others, and the optimal model architecture and data representation for chemical generation is still an area of active research. These models have been extensively applied toward the design of drug-like molecules, but other applications also include the generation of novel solar cell materials [21] and crystalline species [22].

Several of the early papers on generative chemical models have noted low generative validity [10] and diversity [23] when using typical chemical representations (e.g. using SMILES, InChI, and grammar based representations) and architectures. Thus, intense research has been devoted to developing model architectures that guarantee or improve the chemical validity and uniqueness of the molecules produced by deep generative models [24–29]. In contrast, the role of training data and the possible impact of including additional chemical property information during training remains largely unexplored. In particular, although generative models are typically trained using abundant syntactic data (*i.e.* molecules with valid Lewis structures from databases like QM9 [30] and ZINC [31]) limited chemical property data is typically incorporated into model training. Likewise, research on which properties are conducive to inverse design has been scarce, and the properties that have been investigated tend to be either narrow in scope, such as simple cheminformatics data like molecular weight and the water-octanol partition coefficient, or properties that may not be directly verified and are instead based on similarity to known lead molecules [12]. A general approach for targeted generation of compounds based on general quantum chemical properties has not yet been developed.

In the present work we investigate whether transfer learning [32–37] (TL) improves the performance of generative models in comparison with baseline models trained to generate molecules with a single property. We focus on multi-task TL [38–41], whereby modifications to model architecture are minimal, save only through the inclusion of additional property prediction tasks. Since multi-task TL is largely model independent, it provides a potentially complementary strategy for improving generative models to contemporary efforts to modify model architecture and data representation. The hypothesis driving our exploration of multi-task TL for generative models is that the limited validity and uniqueness exhibited by many generative models is a symptom of insufficient chemical property data being utilized during training. In particular, elementary chemical properties including internal energy, bandgap, and zero-point vibrational energy (ZPVE) can be routinely calculated from quantum chemistry and may provide relevant information about validity and molecular stability that cannot be learned from molecular structures alone. This hypothesis is explored using a generative variational autoencoder model with multi-task TL trained on the QM9 dataset, along with semi-empirical calculations for validation of generative predictions. We consider internal energy at zero Kelvins (U_0), zero-point vibrational energy (ZPVE), and HOMO-LUMO gap (E_g) evaluated at the DFT level as multi-task prediction properties for training, and compare how generative performance with respect to validity, uniqueness, and property specificity is affected by inclusion of additional ancillary chemical property data. Since QM9 is a database of small molecules, relatively few high $|U_0|$ samples or low E_g samples exist in the training data. Thus, by characterizing generative results in these property spaces we can characterize how generative performance in an extrapolative task is affected by TL. For targeted structure generation, we observe that multi-task TL increases the percentage of valid and unique high $|U_0|$ structures up to sevenfold with increased property specificity. Searching within areas of physically accessible property values tends to greatly increase the proportion of valid chemical species generated. Even within property regimes with limited representation in the training data, such as structures within the optical bandgap of 1.5–2.0 eV, the inclusion of additional property data improves the ability of the generative models to discover new valid structures. Thus, multi-task TL can be utilized in both extrapolative and interpolative applications to generate novel compounds with optimized or targeted molecular properties with a higher success rate.

2. Computational methods

2.1. Datasets

All models were trained and evaluated using structures from the QM9 dataset. This dataset contains molecular properties computed at the B3LYP/6-31G(2df,p) level for small molecules ranging from 1 to 9 heavy atoms (C, O, N, F). 80% of the structures and their associated properties were utilized for model training, 10% were withheld as part of a validation set to gauge model performance during training, and the final 10% were kept as an independent testing set for the property prediction tasks [42]. The internal energy at 0 K (U_0), ZPVE, and HOMO-LUMO gap (E_g) were used individually, or in combination, for training multi-task models.

3. ML architecture

Approximately 100 models were evaluated in total, with models distinguished by the number and kind of prediction tasks utilized during training. Unless otherwise noted, model types were evaluated in an ensemble fashion, with ten models trained per ensemble. Sampling results for each model type are averaged across the ensemble to provide uncertainty estimates. The autoencoder model utilized in this study was based on the grammar variational autoencoder developed by Kusner *et al.* The major alteration from the original implementation lies in the use of linear predictor networks during model training and evaluation, which was previously shown to be advantageous for multi-task training [39]. Models were trained using the RMSprop algorithm with a learning rate of 0.001 for 100 epochs. The loss corresponding to predictor mean squared error (MSE) was scaled by 100, Kullback-Leibler (KL) divergence was scaled by 750, and the categorical cross entropy loss weight from encoding/decoding was scaled by 50 initially, before decaying to 1 according to a sigmoid function. These loss weights were selected to ensure well-balanced performance with respect to both reconstruction on the training data and property prediction accuracy. Exemplary trained models, model schematic, and full training details may be found in the Supporting Information.

Over 6 million compounds were cumulatively generated in the present study. Since it is infeasible to characterize all of these compounds at the B3LYP/6-31G(2df,p) level, GFN2-xTB (xTB) was utilized to evaluate U_0 and E_g of newly generated species [43]. DFT level predictions for E_g were estimated from a random forest model trained to predict the difference between xTB and B3LYP/6-31G(2df,p) (figure S1 (stacks.iop.org/MLST/1/045010/mmedia)). The random forest model consists of 120 regressors expanded to max depth and trained to predict the difference in xTB and DFT computed bandgap given a Morgan fingerprint of depth two with 1024 bits. Fingerprints were calculated using the RDKit implementation, and the random forest model followed the Scikit-Learn implementation. xTB predictions for $|U_0|$ were found to be linear correlates for the DFT values (figure S2) and were used without modification.

3.1. Sampling paradigms

In the variational chemical autoencoder studied here, new compounds are generated by sampling points from the latent chemical distribution and passing these points to the decoder. By including an ancillary property prediction task, molecular properties vary linearly along particular directions within the latent space (*i.e.* a principle component). The regions of the latent space to target can thus be determined by linear regression of the target property of training compounds with respect to their position along the principal components. The generative performance of the models was investigated using both extrapolative and interpolative sampling approaches. Extrapolative property sampling was achieved using a one-sided distribution with mean equal to the extreme latent position along the sampled principle component in the training data and variance determined from the variance of the training data along the sampled principle component. All other dimensions were normally sampled according to the training set distribution.

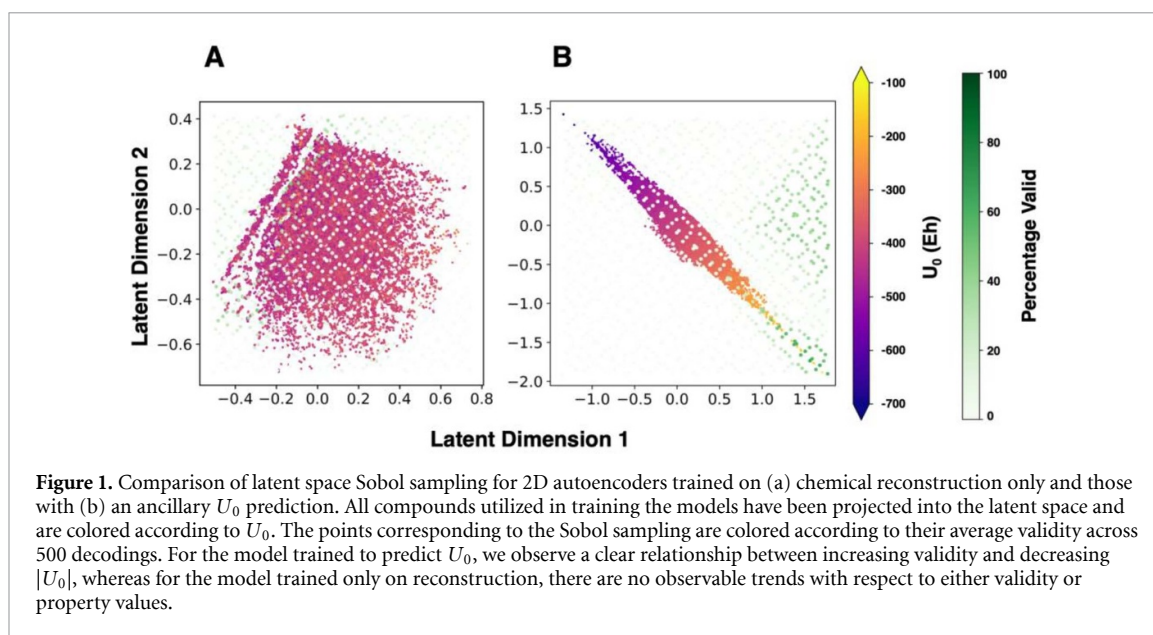
In interpolative sampling, the goal is to recover compounds whose properties lie within a given range (with at least some representation within the training data). The maximum and minimum position along the latent dimension corresponding to the high and low end of the targeted property range were determined via linear regression. This dimension was then sampled from a uniform distribution with specified high and low values. All other dimensions were sampled normally according to the training set distribution.

4. Results and discussion

4.1. Extrapolative sampling

To explore whether adding property prediction tasks can improve generative models we investigated two simplified models: one trained only on encoding and decoding, and another also trained to predict U_0 from the latent encoding. The latent dimensionality was set at two to allow comprehensive sampling to evaluate the validity of generated compounds and for direct visualization of the latent space. One thousand points were sampled from this 2D latent space using the quasi-random Sobol algorithm to maximize the area of the latent space sampled. Due to the stochastic nature of the decoding process, each point was decoded 500 times to determine the average sampling validity. These models are not optimized for either predictive or generative performance, but have been aggressively simplified to illustrate the basic TL mechanism explored in this work. The results from these experiments are presented in figure 1.

This example illustrates two salient features of adding a property prediction task to the generative model. First, by learning to predict a chemical property the models also learn to discriminate between valid and invalid structures. For the autoencoder trained without property prediction tasks (figure 1(a)), the latent space is organized syntactically, with a region containing almost exclusively aromatic compounds in the top left, separated by a gap from other non-aromatic ring-containing structures (figure S3). However, we observe



no discernable trend with respect to decoding validity or $|U_0|$ across the latent dimensions. In contrast, the model trained to predict $|U_0|$ (figure 1(b)) exhibits sampling validity trends that follow the physical property. Second, adding a property prediction task exposes both training data limitations and features of chemistry that affect generative behavior. The sampling validity of the model trained to predict $|U_0|$ approaches nearly 100% for small molecules (*i.e.* low $|U_0|$) and lower validity for large structures (*i.e.* high $|U_0|$). This is consistent with the observation that small molecules are easier to validly decode and are well represented in the training data. In contrast, high $|U_0|$ structures are relatively rare in QM9 given that the dataset is curated for small molecules.

Motivated by this example, we then explored whether generative validity also follows chemical property trends in higher dimensional autoencoders that are more typical of generative applications [9, 10, 18, 44]. In particular, higher dimensional latent spaces are required to effectively compress chemical data when multiple properties are utilized during training (figures S4 and S5). We compared the generative performance of two 56-dimensional autoencoders, one trained only on encoding and decoding and the other also trained to predict U_0 from the latent encoding. Both models were sampled in an extrapolative mode along their first principal component as determined by the latent encodings of the training data, with 30 000 unique structures generated for each model type. For the model trained without property prediction, the extrapolative sampling validity is found to be $25\% \pm 7\%$ with uniqueness of $23\% \pm 5\%$, regardless of direction along the first principal component and consistent with no organization of the latent space with respect to validity. Conversely, sampling in the low absolute internal energy regime of the jointly trained model nearly doubles the average validity, raising it to $45\% \pm 3\%$, but drastically reduces the uniqueness to $3\% \pm 2\%$, as there are only a small number of C,N,H,O, and F containing molecules as U_0 approaches zero. Conversely, extrapolating to high $|U_0|$ produces a valid structure in only $3\% \pm 2\%$ of trials with similar values for uniqueness. These results establish that validity trends in high-dimensional autoencoders follow chemical property values, consistent with the experiments on the low-dimensional models.

It is clear that the inclusion of chemical property data during generative model training affects sampling validity and suggests that supplying additional chemical information may further improve generative performance. We therefore examined the effect of including additional property prediction tasks based on available QM9 data, including ZPVE, which is expected to scale with the number of bonds in the molecule, and E_g , which is not anticipated to correlate with the internal energy of a compound, and repeated the U_0 extrapolation outlined above. The validity and uniqueness statistics for the baseline model trained only on U_0 , and models trained on U_0 /ZPVE, and U_0 /ZPVE/ E_g are summarized in figure 2. The addition of ancillary ZPVE and E_g prediction tasks improves the ability of the multitask models to generate high absolute internal energy structures. Inclusion of ZPVE alone leads to an increase in the average sampling validity and fraction of unique structures. Including an additional E_g prediction task results in little further change in sampling validity, but the proportion of unique structures generated increases and the variance of both quantities decrease.

To investigate if the increased validity and uniqueness of the multi-task models is also reflected in increased property specificity, we have histogrammed the predicted U_0 of the generated compounds for each

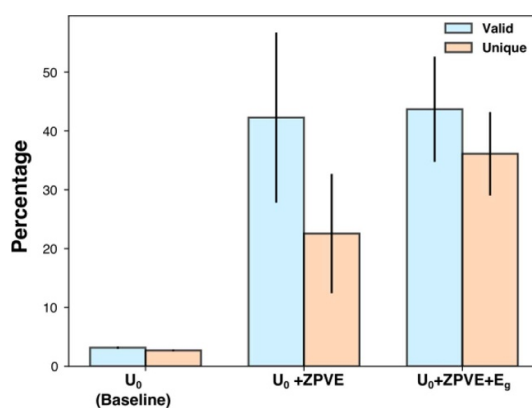


Figure 2. Average sampling validity and uniqueness for the high $|U_0|$ extrapolation trial. Results are averaged across ten distinct models for each training paradigm, with error bars denoting standard deviation. Baseline models exhibit very limited ability to generate structures with high $|U_0|$. The addition of ancillary property prediction tasks greatly improves the generative utility of these models.




Figure 3. Distribution of U_0 for structures for the training set (a) compared with structures generated from extrapolation along the first principal component of models trained on (b) encoding and decoding alone, (c) an ancillary U_0 prediction task, (d) U_0 and $ZPVE$ prediction tasks, and (e) U_0 , $ZPVE$, and E_g prediction tasks. For each paradigm, 3000 unique structures are generated across the ten duplicate models for a total of 30 000 structures. The mean of each distribution is denoted with a dashed red line and the largest $|U_0|$ value in the training data is indicated by the dashed green line. The percentage of generated compounds with $|U_0|$ greater than observed in the training data is shown on the right.

model in figure 3. We note a root mean square deviation of ~ 11 Hartrees between U_0 calculated at the DFT and xTB level (figure S2), thus we expect this to be the minimum uncertainty in the resulting estimates and also contributes to width of the resulting distributions. Extrapolative sampling of the base autoencoder (figure 3(b)) exhibits a broader distribution of U_0 values than observed in the training distribution, however the mean U_0 is consistent with the training data and $<5\%$ of the structures exceed the maximum U_0 of the training data (figure 3(a)). Figure 3(c) shows that including the U_0 prediction task leads to an increase in the number of high $|U_0|$ structures and shifts the mean relative to the training data; however, the mean of the predicted distribution is still within the training distribution, indicating a limited ability to extrapolate beyond the training data. Figure 3(d) shows that inclusion of the $ZPVE$ prediction task shifts the distribution toward higher $|U_0|$ structures well beyond the training data, with $\sim 20\%$ of the generated structures displaying $|U_0|$ greater than the maximum value in the training set. Thus, the increase in valid/unique structures observed in figure 2 comes from the higher density of desired structures within the sampling region. Interestingly, while the further addition of a E_g prediction task in figure 4(e) increases the validity and uniqueness of the generated structures, the bias toward higher $|U_0|$ structures has been removed. It appears that inclusion of E_g saturates the latent space with valid structures, albeit biased toward the training population. The disparate impact on U_0 extrapolation of including the $ZPVE$ and E_g prediction tasks can be understood in terms of their correlation with U_0 . $ZPVE$ scales with molecular size, and so supplying this chemical information improves the extrapolative sampling of high $|U_0|$ structures; however, E_g is weakly correlated with U_0 , but still provides chemical information that improves the overall validity of sampled molecules.

4.2. Interpolative sampling

The extrapolative sampling that was investigated above is representative of chemical discovery applications where champion property values are being sought for new compounds outside of the convex hull of the




Figure 4. Distribution of E_g for the training data (a) and structures generated from models trained to predict E_g by targeting (b) 1.5–2.0 eV, (c) 5.5–6.0 eV, and (d) 9.5–10.0 eV. While (c) and (d) show good specificity, the model is unable to resolve structures in the 1.5–2.0 eV range (b). For each target, 3000 unique structures are generated across the ten duplicate models for a total of 30 000 structures. The median of each distribution is indicated by a dashed red line. Targeted regions are highlighted in blue.

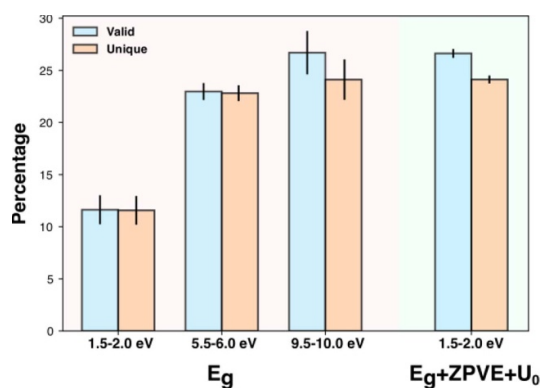
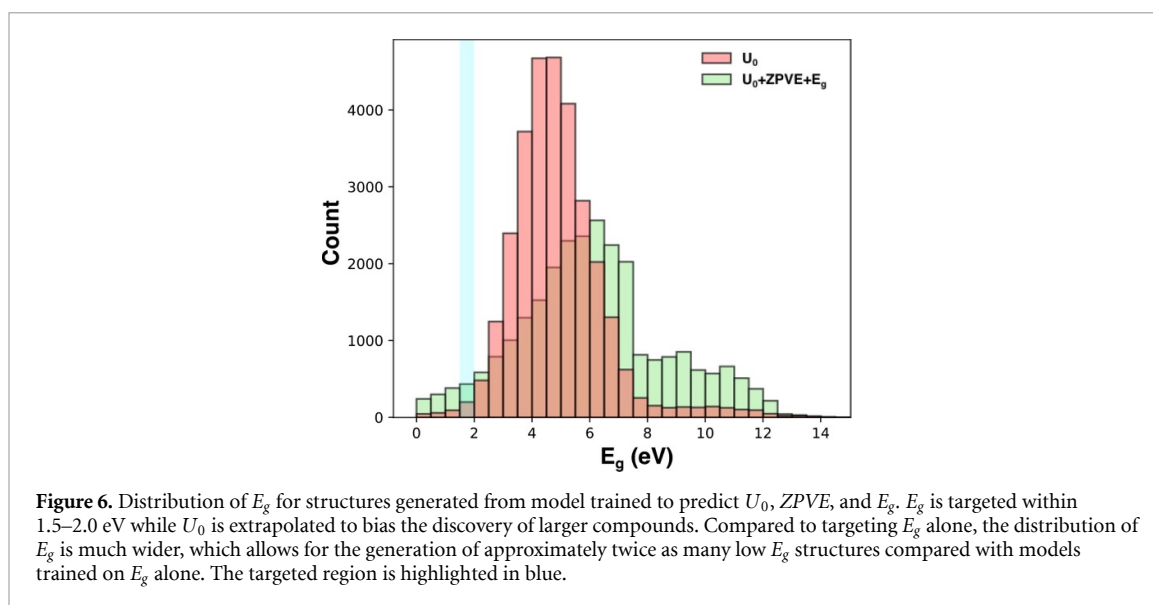


Figure 5. Average sampling validity and uniqueness for the three targeted E_g paradigms. Results are averaged across ten distinct models for each training paradigm, with error bars denoting standard deviation. The models show difficulty in generating compounds within the poorly represented 1.5–2.0 eV range, and are comparatively much stronger in generating structures with E_g between 5.5–6.0 and 9.5–10.0. The addition of ancillary U_0 and ZPVE prediction tasks greatly increases the proportion of valid and unique structures generated in the 1.5–2.0 eV range.

training data. In contrast, interpolative sampling is relevant to applications where a range of property values are desired with some representation in the training data. For instance, in photovoltaic applications it is relevant to target structures within the optical bandgap of 1.0–2.0 eV rather than compounds with extreme values. To investigate the baseline performance of interpolative sampling we trained autoencoders on an E_g prediction task. Ten separate models were trained to evaluate model variance, and novel compounds were generated for the target E_g ranges of 1.5–2.0, 5.5–6.0, and 9.5–10.0 eV based on the principle component corresponding to E_g (see methods). For each range, 3000 unique structures for each of the ten models were decoded and subjected to xTB calculation to determine E_g .

The histograms of E_g values resulting from the interpolative sampling procedure are shown in figure 4. Although all of the targeted ranges have some representation in the training data, the impact of data imbalance is apparent in the results. The generative performance in the 5.5–6.0 eV range, which is well represented in the training set, exhibits good specificity (figure 4(c)). A narrow distribution develops with its most frequent value within the targeted region. Structures with bandgap between 9.5 and 10.0 eV (figure 4(d)) are much less well represented in the training data, and consequently the distribution is much broader and is not centered about the target range, although generation is still shifted toward high E_g structures and a large number of structures are still recovered within the targeted region. The situation is much starker for the lowest bandgap target of 1.5–2.0 eV (figure 4(b)), which shows almost no representation in the training data (figure 4(a)) due to the rarity of such low bandgap structures in a dataset comprised of small molecules. The models exhibit limited ability to target compounds within this region, as evidenced by the distribution centered at 4.5–5.0 eV. Only 200 structures with a bandgap of 1.5–2.0 were generated by these models, out of a total of 30 000 generated compounds. Sampling in this region also tends to produce a much lower fraction of valid/unique structures than the higher E_g targets (figure 5).



Given the scarcity of training compounds exhibiting E_g in the 1.5–2.0 eV range, this generation task is closely analogous to the extrapolative sampling experiments performed for high $|U_0|$ structures. The amount of structures with bandgap of 1.5–2.0 eV is severely limited within the QM9 database, making this a particularly difficult region to target. Because QM9 exhausts small molecule space, there is also a limited amount of target compounds to actually discover within the space spanned by the training data. In order to effectively target this underrepresented region, it is also necessary to expand the search beyond QM9 toward larger molecules. To investigate if multi-task TL could improve targeted generative performance, we followed the same procedure of high $|U_0|$ extrapolation while also targeting structures with E_g within 1.5–2.0 eV. We trained ten separate models on $E_g/ZPVE/U_0$ prediction tasks and performed targeted sampling for E_g in the 1.5–2.0 eV range while extrapolating along the U_0 principle component to facilitate the generation of larger structures. As shown in figure 5, the inclusion of these ancillary prediction tasks more than triples the amount of unique structures generated. However, it is clear from observing the bandgap histogram in figure 6 that the increase in unique structures is not due to greater specificity. In fact, the distribution has been shifted toward higher E_g structures compared to the model trained to predict E_g alone, and the distribution has broadened. It is this broadening of the distribution that is responsible for the increase in the number of unique structures generated, however it also allows the model to generate compounds with a bandgap of 1.5–2.0 eV. In particular, the number of structures within this range has more than doubled, with 440 structures compared to 200 for models trained on E_g prediction alone. The additional chemical property information provided by the $|U_0|$ and $ZPVE$ prediction tasks has been transferred to the task of generating unique structures with targeted properties by improving the ability of the model to sample within these underrepresented regions of chemical space. For the other bandgap targets, the inclusion of $ZPVE$ and U_0 data also serves to broaden the distribution of generated compounds (figure S6); however, as these regions are already targeted effectively, it is not advantageous to add in the additional property data in these cases.

5. Conclusions

These results demonstrate that multi-task TL can be extended beyond property prediction toward improving generative performance. We observe that including chemical property data during generative model training provides complementary information to the syntactic, structural data that is typically used during training, with a generic positive impact on generative validity and performance on extrapolation tasks. In particular, by constraining the search to physically accessible property values, the validity of generated species is increased. For extrapolative sampling, it may prove difficult to generate structures that are not well represented in the training data due to significant differences in atom connectivity and topology; however, the information learned in an ancillary property prediction task can be transferred to the generation task, giving the model enough additional information to successfully resolve these new structures. For targeted structure generation within property regimes that may not be well represented in the training data, this mechanism may also be exploited to help the model effectively sample these underrepresented regions of chemical space and resolve compounds with the desired property values. The choice of properties utilized in this study was largely motivated by their availability and ease of calculation, however we anticipate that this

effect can be employed in other extrapolative applications to generate novel compounds with optimized molecular properties with a higher success rate.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. DGE-1333468. Work performed by B.M.S. was made possible through the Air Force Office of Scientific Research (AFOSR) under support provided by the Organic Materials Chemistry Program (Grant No. : FA9550-18-S-0003, Program Manager: Dr Kenneth Caster). This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation Grant No. ACI-1548562. Simulations were performed on the Comet supercomputer at the University of California, San Diego, under the Allocation no. TG-CHE190014.

Data availability statement

All data that support the findings of this study are included within the article (and any supplementary information files).

ORCID iD

Brett M Savoie  <https://orcid.org/0000-0002-7039-4039>

References

- [1] Radford A, Metz L and Chintala S 2016 Unsupervised representation learning with deep convolutional generative adversarial networks *ArXiv151106434 Cs*
- [2] Briot J-P, Hadjeres G and Pachet F-D 2019 Deep learning techniques for music generation – a survey. *ArXiv170901620 Cs*
- [3] Sanchez-Lengeling B and Aspuru-Guzik A 2018 Inverse molecular design using machine learning: generative models for matter engineering *Science* **361** 360–5
- [4] Schwalbe-Koda D and Gómez-Bombarelli R 2020 Generative models for automatic chemical design *Machine Learning Meets Quantum Physics* ed K T Schütt, S Chmiela, O A von Lilienfeld, A Tkatchenko, K Tsuda and K-R Müller (Cham: Lecture Notes in Physics Springer International Publishing) pp 445–67
- [5] Butler K T, Davies D W, Cartwright H, Isayev O and Walsh A 2018 Machine learning for molecular and materials science *Nature* **559** 547–55
- [6] Dimitrov T, Kreisbeck C, Becker J S, Aspuru-Guzik A and Saikin S K 2019 Autonomous molecular design: then and now *ACS Appl. Mater. Interfaces* **11** 24825–36
- [7] Schmidt J, Marques M R G, Botti S and Marques M A L 2019 Recent advances and applications of machine learning in solid-state materials science *Npj Comput. Mater.* **5** 1–36
- [8] Liu Y, Zhao T, Ju W and Shi S 2017 Materials discovery and design using machine learning *J. Materiomics* **3** 159–77
- [9] Blaschke T, Olivecrona M, Engkvist O, Bajorath J and Chen H 2018 Application of generative autoencoder in de novo molecular design *Mol. Inf.* **37** 1700123
- [10] 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 and Aspuru-Guzik A 2018 Automatic chemical design using a data-driven continuous representation of molecules *ACS Cent. Sci.* **4** 268–76
- [11] Jin W, Barzilay R and Jaakkola T 2019 Junction tree variational autoencoder for molecular graph generation. *ArXiv180204364 Cs Stat*
- [12] Gupta A, Müller A T, Huisman B J H, Fuchs J A, Schneider P and Schneider G 2018 Generative recurrent networks for de novo drug design *Mol. Inf.* **37** 1700111
- [13] Segler M H S, Kogej T, Tyrchan C and Waller M P 2018 Generating focused molecule libraries for drug discovery with recurrent neural networks *ACS Cent. Sci.* **4** 120–31
- [14] Arús-Pous J, Blaschke T, Ulander S, Reymond J-L, Chen H and Engkvist O 2019 Exploring the GDB-13 chemical space using deep generative models *J. Cheminf.* **11** 20
- [15] Guimaraes G L, Sanchez-Lengeling B, Outeiral C, Farias P L C and Aspuru-Guzik A. 2017 Objective-reinforced generative adversarial networks (ORGAN) for sequence generation models. *ArXiv170510843 Cs Stat*
- [16] Sanchez-Lengeling B, Outeiral C, Guimaraes G L and Aspuru-Guzik A. 2017 Optimizing distributions over molecular space an objective-reinforced generative adversarial network for inverse-design chemistry (ORGANIC) *figshare*
- [17] De Cao N and Kipf T. 2018 MolGAN: an implicit generative model for small molecular graphs *ArXiv180511973 Cs Stat*
- [18] Kadurin A, Aliper A, Kazennov A, Mamoshina P, Vanhaelen Q, Khrabrov K and Zhavoronkov A 2016 The cornucopia of meaningful leads: applying deep adversarial autoencoders for new molecule development in oncology *Oncotarget* **8** 10883–90
- [19] Kadurin A, Nikolenko S, Khrabrov K, Aliper A and Zhavoronkov A 2017 DruGAN: an advanced generative adversarial autoencoder model for de novo generation of new molecules with desired molecular properties in silico *Mol. Pharmaceutics* **14** 3098–104
- [20] Polykovskiy D, Zhebrak A, Vetrov D, Ivanenkov Y, Aladinskiy V, Mamoshina P, Bozdaganyan M, Aliper A, Zhavoronkov A and Kadurin A 2018 Entangled conditional adversarial autoencoder for de novo drug discovery *Mol. Pharmaceutics* **15** 4398–405
- [21] Jørgensen P B, Mesta M, Shil S, García Lastra J M, Jacobsen K W, Thygesen K S and Schmidt M N 2018 Machine learning-based screening of complex molecules for polymer solar cells *J. Chem. Phys.* **148** 241735
- [22] Noh J, Kim J, Stein H S, Sanchez-Lengeling B, Gregoire J M, Aspuru-Guzik A and Jung Y 2019 Inverse design of solid-state materials via a continuous representation *Matter* **1** 1370–84

- [23] Behrend M. 2017 ChemGAN challenge for drug discovery: can AI reproduce natural chemical diversity? *ArXiv170808227 Cs Stat*
- [24] Li Y, Vinyals O, Dyer C, Pascanu R and Battaglia P 2018 Learning deep generative models of graphs *ArXiv180303324 Cs Stat*
- [25] Simonovsky M and Komodakis N 2018 GraphVAE: towards generation of small graphs using variational autoencoders *ArXiv180203480 Cs*
- [26] Griffiths -R-R and Miguel Hernández-Lobato J 2020 Constrained Bayesian optimization for automatic chemical design using variational autoencoders *Chem. Sci.* **11** 577–86
- [27] Bradshaw J, Paige B, Kusner M J, Segler M and Hernández-Lobato J M A 2019 A model to search for synthesizable molecules *Advances in Neural Information Processing Systems* 7937–7949
- [28] Janet J P, Chan L and Kulik H J 2018 Accelerating chemical discovery with machine learning: simulated evolution of spin crossover complexes with an artificial neural network *J. Phys. Chem. Lett.* **9** 1064–71
- [29] Winter R, Montanari F, Noé F and Clevert D-A 2019 Learning continuous and data-driven molecular descriptors by translating equivalent chemical representations *Chem. Sci.* **10** 1692–701
- [30] Ramakrishnan R, Dral P O, Rupp M and von Lilienfeld O A 2014 Quantum chemistry structures and properties of 134 kilo molecules *Sci. Data* **1** 140022
- [31] Sterling T and Irwin J J 2015 ZINC 15 – ligand discovery for everyone *J. Chem. Inf. Model.* **55** 2324–37
- [32] Pan S J and Yang Q A 2010 Survey on transfer learning *IEEE Trans. Knowl. Data Eng.* **22** 1345–59
- [33] Weiss K, Khoshgoftaar T M and Wang D A 2016 Survey of transfer learning *J. Big Data* **3** 9
- [34] Grambow C A, Li Y-P and Green W H 2019 Accurate thermochemistry with small data sets: a bond additivity correction and transfer learning approach *J. Phys. Chem. A* **123** 5826–35
- [35] Smith J S, Nebgen B T, Zubatyuk R, Lubbers N, Devereux C, Barros K, Tretiak S, Isayev O and Roitberg A E 2019 Approaching coupled cluster accuracy with a general-purpose neural network potential through transfer learning *Nat. Commun.* **10** 1–8
- [36] Sultan M M and Pande V S 2018 Transfer learning from Markov models leads to efficient sampling of related systems *J. Phys. Chem. B* **122** 5291–9
- [37] Zeng W-F, Zhou -X-X, Zhou W-J, Chi H, Zhan J and He S-M 2019 MS/MS spectrum prediction for modified peptides using pdeep2 trained by transfer learning *Anal. Chem.* **91** 9724–31
- [38] Hutchinson M L, Antono E, Gibbons B M, Paradiso S, Ling J and Meredig B 2017 Overcoming data scarcity with transfer learning. *arXiv:1711.05099*
- [39] Iovanac N C and Savoie B M 2020 Simpler is better: how linear prediction tasks improve transfer learning in chemical autoencoders *J. Phys. Chem. A* **124** 3679–85
- [40] Ramsundar B, Kearnes S, Riley P, Webster D, Konerding D and Pande V 2015 Massively multitask networks for drug discovery. *arXiv:1502.02072*
- [41] Iovanac N C and Savoie B M 2019 Improved chemical prediction from scarce data sets via latent space enrichment *J. Phys. Chem. A* **123** 4295–302
- [42] Wu Z, Ramsundar B, Feinberg N, Gomes E, Geniesse J, Pappu C ; S, Leswing A, Pande K and MoleculeNet: V 2018 A benchmark for molecular machine learning *Chem. Sci.* **9** 513–30
- [43] Bannwarth C, Ehlert S and Grimme S 2019 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.* **15** 1652–71
- [44] Jørgensen P B, Schmidt M N and Winther O 2018 Deep generative models for molecular science *Mol. Inf.* **37** 1700133