
Evaluating the diversity and utility of materials proposed by generative models

Alexander New¹ Michael J. Pekala¹ Elizabeth A. Pogue¹ Nam Q. Le¹ Janna Domenico¹ Christine D. Piatko¹
Christopher D. Stiles¹

Abstract

Generative machine learning (ML) models can use data generated by scientific modeling to create large quantities of novel material structures. Here, we assess how one state-of-the-art generative model, the physics-guided crystal generation model (PGCGM), can be used as part of the inverse design process. We show that the default PGCGM’s input space is not smooth with respect to parameter variation, making material optimization difficult and limited. We also demonstrate that most generated structures are predicted to be thermodynamically unstable by a separate property-prediction model, partially due to out-of-domain data challenges. Our findings suggest how generative models might be improved to enable better inverse design.

1. Introduction

Inverse design—the discovery of materials with targeted properties—remains an important task in materials science (Wang et al., 2022). In recent years, progress has been made, spurred by the use of machine learning (ML) and iterative design loops (Pogue et al., 2022; Goodall et al., 2022; Stanev et al., 2018; Attia et al., 2020; Zhang et al., 2020; Baird et al., 2022; Wines et al., 2023), but further challenges remain. A key reason is the lack of high-quality experimental and computational data, compared to similar fields like chemistry (Xie et al., 2022).

Computational and experimental characterization have enabled the generation of scientific materials databases like Materials Project (MP) (Jain et al., 2013), Open Quantum Materials Database (OQMD) (Saal et al., 2013; Kirklin et al., 2015), and Inorganic Crystal Structure Database

¹Research and Exploratory Development Department, Johns Hopkins University Applied Physics Laboratory, Laurel, Maryland, USA. Correspondence to: Alexander New <alex.new@jhuapl.edu>.

Accepted after peer-review at the 1st workshop on Synergy of Scientific and Machine Learning Modeling, SynS & ML ICML, Honolulu, Hawaii, USA. July, 2023. Copyright 2023 by JHU/APL.

(ICSD) (Belsky et al., 2002), which contain hundreds of thousands of material structures. However, these still sample only a small fraction of possible materials. To enable the discovery of additional materials structures, these scientific databases are used as training data for generative ML models like generative adversarial networks (GANs) (Goodfellow et al., 2014) and variational autoencoders (VAEs) (Kingma & Welling, 2014). Via additional property-prediction models, generated materials can be assessed for suitability in design tasks. Identified materials can be studied in detail, synthesized, and characterized. In recent years, a number of different ML models for material generation have been developed (Zhao et al., 2021; Kim & Dordevic, 2022; Ren et al., 2022; Court et al., 2020; Zhao et al., 2023; Alverson et al., 2023), and ML has been successfully used for property prediction (Goodall & Lee, 2020; New et al., 2022; Xie & Grossman, 2018; Sanyal et al., 2018; Chen et al., 2019; Park & Wolverton, 2020; Choudhary & DeCost, 2021).

In this work, we evaluate a state-of-the-art generative model, the physics-guided crystal generation model (PGCGM) (Zhao et al., 2023) used to generate diverse and usable material structures. In Section 2, we detail the generation process. In Section 3, we identify a potential issue with the use of PGCGM for inverse optimization—namely, a lack of smoothness with respect to property and validity variation in the model’s input space. In Section 4, we use a second ML model to assess the stability of generated materials and highlight a lack of diversity in structures. We conclude with recommendations for how to enhance generative models for use in inverse design of materials.

2. Generating novel material structures

Our study uses a state-of-the-art generative model, PGCGM (Zhao et al., 2023). The PGCGM extends the Wasserstein GAN (Arjovsky et al., 2017) to ternary material generation. Given three constituent atoms and space group of a ternary material system, PGCGM predicts possible atom coordinates and lattice parameters. Generated unit cells are saved into crystallographic information file (CIF) format and can be read by tools like `pymatgen` (Ong et al., 2013) and `jarvis` (Choudhary et al., 2020) for use and analysis. PGCGM was shown to perform well in terms of

material validity, diversity, and property similarity to training data, compared to other models like CubicGAN (Zhao et al., 2021) and the Fourier-transformed crystal properties (FTCP) (Ren et al., 2022) method. PGCGM trains a generator f_θ that satisfies:

$$B, p = f_\theta(z, E, s), \quad (1)$$

where the outputs are: $B = (b_0, b_1, b_2) \in \mathbb{R}^{3 \times 3}$, the atom fractional coordinates for each constituent element; and $p = (a, b, c, \alpha, \beta, \gamma)$, the structure’s lattice parameters; and the inputs are: $z \in \mathbb{R}^{128}$, a vector with entries sampled from a standard normal distribution; E , the three constituent atoms; and s , a space group. Thus, B and p can be used to construct a generated material’s CIF.

We use the pretrained PGCGM model and post-processing scripts available on GitHub¹ and generate two sets of material structures. In the first, we fix the constituent atoms and the space group and generate 5,000 samples from f_θ for the Ta–Ge–As system with the $Pm\bar{3}m$ space group; of these, 241 are valid material structures². We analyze how the GAN sample z impacts material structure of these materials in Section 3. In the second set, we sample sets of constituent atoms and space groups and generate 500,000 samples from f_θ ; of these, 27,116 are valid. The validity rate ($27,116/500,000 \approx 5.4\%$) for this sample is slightly lower but comparable to the 7.14% validity rate (Zhao et al., 2023) report.

We analyze the predicted stability of these structures in Section 4. Figure 6 and Figure 7 in Appendix A show example structures generated by PGCGM.

3. Assessing smoothness of GAN input space

In the materials design problem, we look for a material that maximizes a suitability or property function Outcome (e.g., critical temperature for superconductors or solubility for drug-like molecules). Using PGCGM, this problem can be formulated as:

$$\hat{B}, \hat{p} = \arg \max_{z, E, s} \text{Outcome}(B(z, E, s), p(z, E, s)), \quad (2)$$

where B, P, z, S, E are defined in Eq. 1. Given a differentiable suitability function Outcome and fixed space group s and elements E , Eq. 2 suggests the use of gradient-based methods to optimize Outcome by differentiating it with respect to z . Similar approaches are often used with VAEs for inverse design (Gómez-Bombarelli et al., 2018; Xie et al., 2022) and with GANs for image editing (Zhu et al., 2020).

¹<https://github.com/MilesZhao/PGCGM>

²Following the post-processing scripts of (Zhao et al., 2023), by “valid”, we mean that their estimated space group (from pymatgen) matches their input space group s , and atoms of the same type that occupy approximately the same spatial location can be merged together.




Figure 1. We use UMAP to project PGCGM inputs (z in Eq. 1) corresponding to valid CIFs from the Ta–Ge–As system into two dimensions. Nearby points often have different property values, suggesting that the PGCGM input-space is not generally smooth with respect to variation in material properties, making PGCGM difficult to use as part of gradient-based material optimization. Statistical analyses (Section 3) further support this observation.

Effective optimization requires two conditions be met: (1) small perturbations to a PGCGM input z corresponding to a valid (invalid) CIF should continue to produce a valid (invalid) CIF, and (2) perturbing a PGCGM input z corresponding to a valid CIF should not drastically change properties of the unit cell (e.g., the gradient of Outcome should be Lipschitz-continuous with respect to z , which can enable convergence of gradient descent techniques in combination with other regularity conditions (Carmon et al., 2020; Arjevani et al., 2023)).

3.1. Global smoothness analysis

Using the 241 PGCGM inputs z corresponding to valid structures, we find that the pretrained PGCGM’s input space does not appear to be globally smooth. In Figure 1, we take the 241 PGCGM inputs z corresponding to valid structures and project them into two dimensions using uniform manifold approximation and projection (UMAP) (McInnes et al., 2020). It is common for a point and its nearest neighbor in the embedding to have very different property values.

We also conduct a more extensive analysis on the relationship between the PGCGM input space and the set of samples $z \in \mathbb{R}^{128}$ corresponding to valid materials. First, we use the Kolmogorov-Smirnov (KS) test to see if the distribution of each component of the samples z was distinct from the standard normal distribution that generated all of the samples and adjust the 128 p -values with the Benjamini-Hochberg (Benjamini & Hochberg, 1995) scheme. Four z components (out of 128) were found to be significantly distinct from the standard normal distribution, suggesting that z samples corresponding to valid structures and invalid structures are often statistically similar, which is evidence against the first condition. We also train random forest models to predict unit cell density and volume given z ; the density-prediction model had an out-of-bag R^2 estimate of 0.16, and the volume-prediction model had an out-of-bag R^2 estimate of 0.20. This suggests that the relationship between PGCGM input z and resultant unit cell’s properties is difficult to model, providing evidence against the second condition. We repeated this analysis with other material systems and saw comparable results.

3.2. Local smoothness analysis

Our conclusions in Section 3.1 relied on a projection of a high-dimensional space to assess smoothness. To provide additional perspective, we perform a perturbation-based local smoothness analysis. This enables us to explore property variation in local neighborhoods of GAN inputs z corresponding to valid material structures. Specifically, we start by sampling a point z_0 from the GAN’s input distribution $\mathcal{N}(0_{128}, I_{128})$. We sample N random directions r_n from the surface of the unit sphere in \mathbb{R}^{128} and construct a uniform discretization $0 < \epsilon_1 < \dots < \epsilon_M$ of the range $(0, \epsilon_M]$. By evaluating the PGCGM at each point $z_0 + \epsilon_m r_n$ and calculating Outcome of the resultant material $f_\theta(z_0 + \epsilon_m r_n, E, s)$, we see Outcome’s behavior across balls of radius ϵ_m .

These calculations let us estimate Lipschitz constants for property functions by evaluating $|\text{Outcome}(f_\theta(z_0 + \epsilon_m r_n, E, s)) - \text{Outcome}(f_\theta(z_0 + \epsilon_{m'} r_{n'}, E, s))|$ vs. $|\epsilon_m r_n - \epsilon_{m'} r_{n'}|$. Although convergence rates for optimization involve Lipschitz constants of function gradients (i.e., $\nabla_z \text{Outcome}(f_\theta(z_0 + \epsilon_m r_n, E, s))$), these are harder to obtain for PGCGM: the PGCGM’s post-processing non-smoothly merges spatially-adjacent atoms of the same type together. Future work could estimate some of these Lipschitz constants using specially-developed techniques with provable guarantees (Jordan & Dimakis, 2020; 2021).

In this study, we select z_0 as one of the 241 stable materials sampled from the Ta–Ge–As system with the $Pm\bar{3}m$ space group used in Section 3.1, we sample $N = 128$ different directions r_n and evaluate perturbations $r_n \epsilon_m$ up to a maximum radius of $\epsilon_M = 10$. In Figure 2, we show

pairwise property and input distances for the resulting valid CIFs. For smaller input distances ($\|\epsilon_m r_n - \epsilon_{m'} r_{n'}\| \lesssim 4$), the PGCGM input space is locally smooth in the sense that small changes in the input mostly lead to small changes in the property. After that, the range in variation of the properties increases significantly, which aligns with the global smoothness analysis of Section 3.1.




Figure 2. Using a sample from the Ta–Ge–As system, we analyze the local smoothness of the PGCGM’s input space by plotting changes in material properties—density (top) and volume (bottom)—with respect to changes in the PGCGM’s input z . Notably, the input space features contours where $\|\epsilon_m r_n - \epsilon_{m'} r_{n'}\|$ is constant, but the property changes.

4. High-throughput thermodynamic stability predictions

Generative models suggest large numbers of potential material structures, and not all generated structures are guaranteed to be thermodynamically stable. Some of these can be further characterized by experimental or computational techniques, but this characterization is too expensive in time



Figure 3. 12,826 of the 33,172 PGCGM training set structures are predicted to have a negative decomposition enthalpy (left) ΔH_d by ALIGNN, indicating they are predicted to be thermodynamically stable. In contrast, all but 195 of the 27,116 structures generated by PGCGM were predicted to have a positive decomposition enthalpy (right). The median predicted decomposition enthalpy is also lower for the training data than the generated structures: 0.02 eV/atom vs. 0.38 eV/atom.

and cost to be performed for all or even most generated materials. Thus, we trained a ML to predict the thermodynamic stability of the PGCGM-generated structures.

In particular, we use atomistic line graph neural network (ALIGNN) (Choudhary & DeCost, 2021) to predict decomposition enthalpy ΔH_d . ALIGNN is trained on a set of ternary structures and computationally-predicted ΔH_d taken from MP and collected by Bartel et al. 2020. A structure is predicted to be thermodynamically stable if its decomposition enthalpy is negative.

The ALIGNN forward pass is:

$$\Delta H_d = \ell(h) = \ell(a(g)), \quad (3)$$

where ΔH_d is the predicted decomposition enthalpy, ℓ is a linear layer, and $h = a(g) \in \mathbb{R}^{256}$ is the embedding of an input graph g produced by edge-gated graph convolutional layers (Bresson & Laurent, 2017). Our ALIGNN model is trained with the default architecture and hyperparameters.

We evaluate ALIGNN on PGCGM’s training data (33,172 structures taken from MP, ICSD, and OQMD) and the 27,116 PGCGM-generated structures, calculating both their predicted decomposition enthalpies ΔH_d and their embeddings h . In Figure 3, we show that only 195 of the 27,116 generated structures are predicted by ALIGNN to be thermodynamically stable. In contrast, 12,826 of the 33,172 training set structures are predicted to be stable.

We propose and evaluate two hypotheses for this result. First, the training data for and generated structures from PGCGM are potentially out-of-domain for the trained ALIGNN model. In Figure 4, we show that there is a strong correlation between the minimum Euclidean distance of a PGCGM-generated structure’s embedding h to the set of embeddings of the ALIGNN training set and the predicted ΔH_d . Thus, only the generated structures most similar to




Figure 4. The minimum distance between each of the 27,117 generated materials’ embeddings (h in Eq. 3) and the ALIGNN training data’s embeddings has a strong positive correlation with the predicted decomposition enthalpy (shading indicates regions where points are concentrated). Thus, only the generated structures most similar to the training data are predicted to be thermodynamically stable.

training data are predicted to be stable. In Figure 9 in Appendix A, we build on this distance vs. prediction for a held-out subset of the ALIGNN training data; we show that the correlation of Figure 4 is not a general characteristic of ALIGNN predictions.

Beyond this distance-based analysis driving ALIGNN’s predictions, we also provide evidence that the structures generated by PGCGM are distributionally different and less varied than those in its training data. In Figure 5, we take the embeddings of the PGCGM training data and the generated structures and use UMAP (McInnes et al., 2020) to project them into two dimensions. The embeddings of the generated structures are compressed into a tight region of the projection, while the embeddings of the training data are spread more widely across the space. Thus, there appears to be a lack of diversity in the generated structures, with respect to how the ALIGNN model predicts their stability. This is similar to the general problem of mode collapse observed in GANs (Saxena & Cao, 2021), and strategies adopted there could mitigate it in future work.

In Figure 8 (in the Appendix), we repeat this analysis using the embeddings of the PGCGM-generated structures and ALIGNN’s training data from MP. As in Figure 5, the embeddings of the generated structures are more compressed together than the embeddings of the training data.

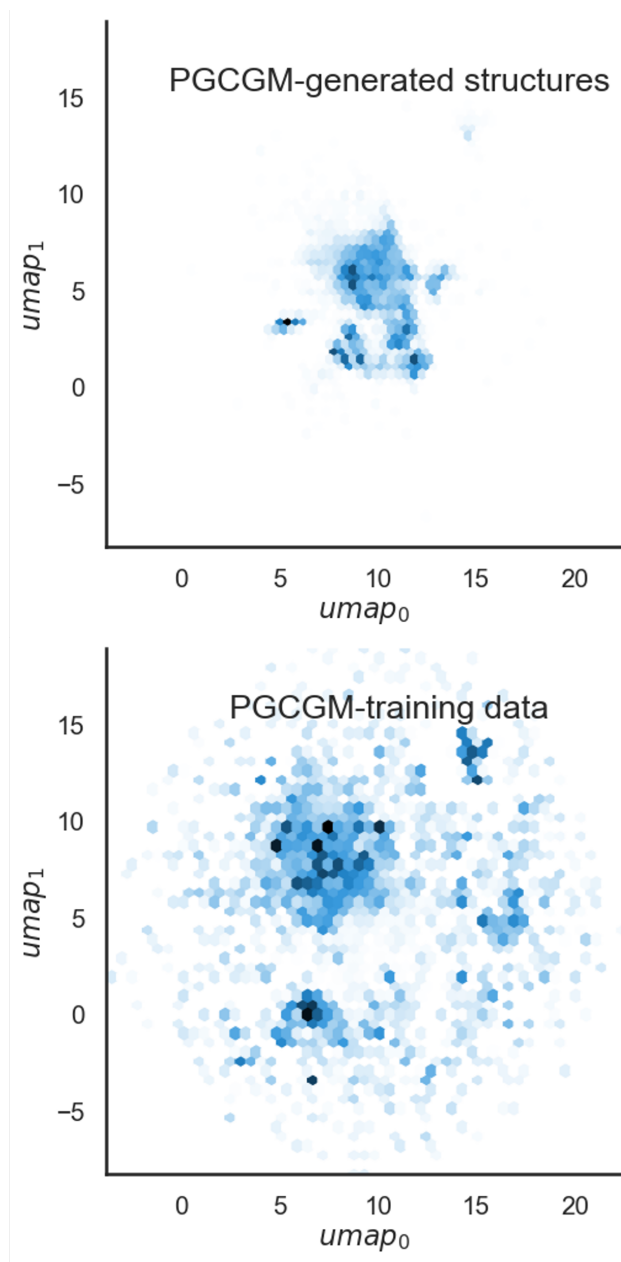


Figure 5. We use UMAP to show that the ALIGNN embeddings of the 27,116 valid CIFs generated by PGCGM are concentrated in a tight region of latent space. In contrast, the ALIGNN embeddings of PGCGM’s training data are spread more widely (shading indicates regions where points are concentrated). This suggests PGCGM’s capacity for generating diverse and varied structures can be further improved. Here, the UMAP embedding was learned from the combination of the PGCGM’s generated structures and training data.

5. Conclusion

Here we assess the ability of generative structure models to assist in designing novel materials. Although current generative models like PGCGM are able to generate large numbers of varied and valid material structures, we show that (1) the PGCGM is not smooth with respect to either material validity or material property variation, making property optimization difficult; and (2) assessing thermodynamic stability of generated structures is difficult due to generated structures being concentrated in tight regions of latent space distant from property-predictor training data. Our particular analyses are not restricted to the PGCGM and can be applied to assess other generative material models such as Kim & Dordevic (2022); Ren et al. (2022); Court et al. (2020); Alverson et al. (2023).

Our findings here suggest potential investigations to fully enable inverse design via generative modeling. For example, the loss formulation of PGCGM or another generative model could be modified to encourage smoothness with respect to property variation in its input space via techniques like optimal transport (Chen et al., 2021). Global smoothness of PGCGM input space is not necessarily needed, as functions that are locally Lipschitz can still be optimized in some settings (Patel & Berahas, 2022).

Although we saw that most PGCGM-generated structures are predicted to be unstable, this may be a consequence of a poor property-prediction model, not a poor input structure. Compared to general neural networks, graph neural networks like ALIGNN have been shown to perform well in certain extrapolation tasks (Xu et al., 2021); however, the distance-to-training set metric we use in Figure 4 does not fully explain if a prediction is accurate (Hirschfeld et al., 2020). Recent work also observes that ALIGNN models can fail specifically in the setting of out-of-domain property prediction for material structures (Li et al., 2023). Challenges have also been identified in the use of graph neural networks to predict multiple material properties (e.g., stability and a design property of interest) (New et al., 2022).

Here we analyze how diverse and stable structures directly predicted by the PGCGM are. A supplementary strategy, taken by (Zhao et al., 2023), is to further refine ML-predicted material structures, either with first-principles calculations, or with techniques like constrained Bayesian optimization (Zuo et al., 2021).

Broader impact

Inverse design of materials is important for many fields of engineering and science, and these fields have robust mechanisms for preventing misuse. Materials discoveries made by ML can thus be beneficial to different aspects of society.

Acknowledgements

This work was supported by internal research and development funding from the Research and Exploratory Development Mission Area of the Johns Hopkins University Applied Physics Laboratory. We also thank the reviewers for their helpful suggestions and comments.

References

- Alverson, M., Baird, S., Murdock, R., Ho, S.-H., Johnson, J., and Sparks, T. Generative adversarial networks and diffusion models in material discovery, 2023. URL <https://doi.org/10.26434/chemrxiv-2022-614pm-v2>.
- Arjevani, Y., Carmon, Y., Duchi, J. C., Foster, D. J., Srebro, N., and Woodworth, B. Lower bounds for non-convex stochastic optimization. *Mathematical Programming*, 199(1):165–214, May 2023. ISSN 1436-4646. doi: 10.1007/s10107-022-01822-7. URL <https://doi.org/10.1007/s10107-022-01822-7>.
- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein GAN, 2017.
- Attia, P. M., Grover, A., Jin, N., Severson, K. A., Markov, T. M., Liao, Y.-H., Chen, M. H., Cheong, B., Perkins, N., Yang, Z., Herring, P. K., Aykol, M., Harris, S. J., Braatz, R. D., Ermon, S., and Chueh, W. C. Closed-loop optimization of fast-charging protocols for batteries with machine learning. *Nature*, 578(7795):397–402, Feb 2020. ISSN 1476-4687. doi: 10.1038/s41586-020-1994-5. URL <https://doi.org/10.1038/s41586-020-1994-5>.
- Baird, S. G., Diep, T. Q., and Sparks, T. D. Discover: a materials discovery screening tool for high performance, unique chemical compositions. *Digital Discovery*, 1:226–240, 2022. doi: 10.1039/D1DD00028D. URL <http://dx.doi.org/10.1039/D1DD00028D>.
- Bartel, C. J., Trewartha, A., Wang, Q., Dunn, A., Jain, A., and Ceder, G. A critical examination of compound stability predictions from machine-learned formation energies. *npj Computational Materials*, 6(1):97, Jul 2020. ISSN 2057-3960. doi: 10.1038/s41524-020-00362-y. URL <https://doi.org/10.1038/s41524-020-00362-y>.
- Belsky, A., Hellenbrandt, M., Karen, V. L., and Luksch, P. New developments in the Inorganic Crystal Structure Database (ICSD): accessibility in support of materials research and design. *Acta Crystallographica Section B*, 58(3 Part 1):364–369, Jun 2002. doi: 10.1107/S0108768102006948. URL <https://doi.org/10.1107/S0108768102006948>.
- Benjamini, Y. and Hochberg, Y. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1):289–300, 2023/05/12/ 1995. ISSN 00359246. URL <http://www.jstor.org/stable/2346101>. Full publication date: 1995.
- Bresson, X. and Laurent, T. Residual gated graph ConvNets. *CoRR*, abs/1711.07553, 2017. URL <http://arxiv.org/abs/1711.07553>.
- Carmon, Y., Duchi, J. C., Hinder, O., and Sidford, A. Lower bounds for finding stationary points i. *Mathematical Programming*, 184(1):71–120, Nov 2020. ISSN 1436-4646. doi: 10.1007/s10107-019-01406-y. URL <https://doi.org/10.1007/s10107-019-01406-y>.
- Chen, B., Bécigneul, G., Ganea, O.-E., Barzilay, R., and Jaakkola, T. Optimal transport graph neural networks, 2021.
- Chen, C., Ye, W., Zuo, Y., Zheng, C., and Ong, S. P. Graph networks as a universal machine learning framework for molecules and crystals. *Chemistry of Materials*, 31(9):3564–3572, May 2019. ISSN 0897-4756. doi: 10.1021/acs.chemmater.9b01294. URL <https://doi.org/10.1021/acs.chemmater.9b01294>.
- Choudhary, K. and DeCost, B. Atomistic line graph neural network for improved materials property predictions. *npj Computational Materials*, 7(1):185, Nov 2021. ISSN 2057-3960. doi: 10.1038/s41524-021-00650-1. URL <https://doi.org/10.1038/s41524-021-00650-1>.
- Choudhary, K., Garrity, K. F., Reid, A. C. E., DeCost, B., Biacchi, A. J., Hight Walker, A. R., Trautt, Z., Hattrick-Simpers, J., Kusne, A. G., Centrone, A., Davydov, A., Jiang, J., Pachter, R., Cheon, G., Reed, E., Agrawal, A., Qian, X., Sharma, V., Zhuang, H., Kalinin, S. V., Sumpter, B. G., Pilia, G., Acar, P., Mandal, S., Haule, K., Vanderbilt, D., Rabe, K., and Tavazza, F. The joint automated repository for various integrated simulations (JARVIS) for data-driven materials design. *npj Computational Materials*, 6(1):173, Nov 2020. ISSN 2057-3960. doi: 10.1038/s41524-020-00440-1. URL <https://doi.org/10.1038/s41524-020-00440-1>.

- Court, C. J., Yildirim, B., Jain, A., and Cole, J. M. 3-d inorganic crystal structure generation and property prediction via representation learning. *Journal of Chemical Information and Modeling*, 60(10):4518–4535, Oct 2020. ISSN 1549-9596. doi: 10.1021/acs.jcim.0c00464. URL <https://doi.org/10.1021/acs.jcim.0c00464>.
- Goodall, R. E., Parackal, A. S., Faber, F. A., Armiento, R., and Lee, A. A. Rapid discovery of stable materials by coordinate-free coarse graining. *Science Advances*, 8(30): eabn4117, 2022.
- Goodall, R. E. A. and Lee, A. A. Predicting materials properties without crystal structure: deep representation learning from stoichiometry. *Nature Communications*, 11(1):6280, Dec 2020. ISSN 2041-1723. doi: 10.1038/s41467-020-19964-7. URL <https://doi.org/10.1038/s41467-020-19964-7>.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, NIPS'14, pp. 2672–2680, Cambridge, MA, USA, 2014. MIT Press.
- 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. Automatic chemical design using a data-driven continuous representation of molecules. *ACS Central Science*, 4(2):268–276, 2018. doi: 10.1021/acscentsci.7b00572. URL <https://doi.org/10.1021/acscentsci.7b00572>. PMID: 29532027.
- Hirschfeld, L., Swanson, K., Yang, K., Barzilay, R., and Coley, C. W. Uncertainty quantification using neural networks for molecular property prediction. *Journal of Chemical Information and Modeling*, 60(8):3770–3780, 2020. doi: 10.1021/acs.jcim.0c00502. URL <https://doi.org/10.1021/acs.jcim.0c00502>. PMID: 32702986.
- Jain, A., Ong, S. P., Hautier, G., Chen, W., Richards, W. D., Dacek, S., Cholia, S., Gunter, D., Skinner, D., Ceder, G., and Persson, K. A. Commentary: The materials project: A materials genome approach to accelerating materials innovation. *APL Materials*, 1:011002, 7 2013. ISSN 2166532X. doi: 10.1063/1.4812323. URL <https://aip.scitation.org/doi/abs/10.1063/1.4812323>.
- Jordan, M. and Dimakis, A. Provable Lipschitz certification for generative models. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 5118–5126. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/jordan21a.html>.
- Jordan, M. and Dimakis, A. G. Exactly computing the local Lipschitz constant of ReLU networks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS'20, Red Hook, NY, USA, 2020. Curran Associates Inc. ISBN 9781713829546.
- Kim, E. and Dordevic, S. V. ScGAN: A generative adversarial network to predict hypothetical superconductors, 2022.
- Kingma, D. P. and Welling, M. Auto-encoding variational Bayes. In Bengio, Y. and LeCun, Y. (eds.), *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. URL <http://arxiv.org/abs/1312.6114>.
- Kirklin, S., Saal, J. E., Meredig, B., Thompson, A., Doak, J. W., Aykol, M., Rühl, S., and Wolverton, C. The Open Quantum Materials Database (OQMD): assessing the accuracy of DFT formation energies. *npj Computational Materials*, 1(1):15010, Dec 2015. ISSN 2057-3960. doi: 10.1038/npjcompumats.2015.10. URL <https://doi.org/10.1038/npjcompumats.2015.10>.
- Li, K., DeCost, B., Choudhary, K., Greenwood, M., and Hattrick-Simpers, J. A critical examination of robustness and generalizability of machine learning prediction of materials properties. *npj Computational Materials*, 9(1):55, Apr 2023. ISSN 2057-3960. doi: 10.1038/s41524-023-01012-9. URL <https://doi.org/10.1038/s41524-023-01012-9>.
- McInnes, L., Healy, J., and Melville, J. UMAP: Uniform manifold approximation and projection for dimension reduction, 2020. URL <https://arxiv.org/abs/1802.03426>.
- New, A., Pekala, M. J., Le, N. Q., Domenico, J., Piatko, C. D., and Stiles, C. D. Curvature-informed multi-task learning for graph networks. In *ICML 2022 2nd AI for Science Workshop*, 2022. URL <https://openreview.net/forum?id=m5RYtApKF0g>.
- Ong, S. P., Richards, W. D., Jain, A., Hautier, G., Kocher, M., Cholia, S., Gunter, D., Chevrier, V. L., Persson, K. A., and Ceder, G. Python materials genomics (pymatgen): A robust, open-source python library for materials analysis. *Computational Materials Science*, 68:314–319, 2013. ISSN 0927-0256. doi: <https://doi.org/10.1016/j.commatsci.2012.10>.

028. URL <https://www.sciencedirect.com/science/article/pii/S0927025612006295>.
- Park, C. W. and Wolverton, C. Developing an improved crystal graph convolutional neural network framework for accelerated materials discovery. *Phys. Rev. Materials*, 4:063801, Jun 2020. doi: 10.1103/PhysRevMaterials.4.063801. URL <https://link.aps.org/doi/10.1103/PhysRevMaterials.4.063801>.
- Patel, V. and Berahas, A. S. Gradient descent in the absence of global lipschitz continuity of the gradients: Convergence, divergence and limitations of its continuous approximation, 2022.
- Pogue, E. A., New, A., McElroy, K., Le, N. Q., Pekala, M. J., McCue, I., Gienger, E., Domenico, J., Hedrick, E., McQueen, T. M., Wilfong, B., Piatko, C. D., Ratto, C. R., Lennon, A., Chung, C., Montalbano, T., Bassen, G., and Stiles, C. D. Closed-loop machine learning for discovery of novel superconductors, 2022.
- Ren, Z., Tian, S. I. P., Noh, J., Oviedo, F., Xing, G., Li, J., Liang, Q., Zhu, R., Aberle, A. G., Sun, S., Wang, X., Liu, Y., Li, Q., Jayavelu, S., Hippalgaonkar, K., Jung, Y., and Buonassisi, T. An invertible crystallographic representation for general inverse design of inorganic crystals with targeted properties. *Matter*, 5(1):314–335, 2022. ISSN 2590-2385. doi: <https://doi.org/10.1016/j.matt.2021.11.032>. URL <https://www.sciencedirect.com/science/article/pii/S2590238521006251>.
- Saal, J. E., Kirklin, S., Aykol, M., Meredig, B., and Wolverton, C. Materials design and discovery with high-throughput density functional theory: The open quantum materials database (OQMD). *JOM*, 65(11): 1501–1509, Nov 2013. ISSN 1543-1851. doi: 10.1007/s11837-013-0755-4. URL <https://doi.org/10.1007/s11837-013-0755-4>.
- Sanyal, S., Balachandran, J., Yadati, N., Kumar, A., Rajagopalan, P., Sanyal, S., and Talukdar, P. MT-CGCNN: Integrating crystal graph convolutional neural network with multitask learning for material property prediction, 2018.
- Saxena, D. and Cao, J. Generative adversarial networks (GANs): Challenges, solutions, and future directions. *ACM Comput. Surv.*, 54(3), may 2021. ISSN 0360-0300. doi: 10.1145/3446374. URL <https://doi.org/10.1145/3446374>.
- Stanev, V., Oses, C., Kusne, A. G., Rodriguez, E., Paglione, J., Curtarolo, S., and Takeuchi, I. Machine learning modeling of superconducting critical temperature. *npj Computational Materials*, 4(1):29, Jun 2018. ISSN 2057-3960. doi: 10.1038/s41524-018-0085-8. URL <https://doi.org/10.1038/s41524-018-0085-8>.
- Wang, J., Wang, Y., and Chen, Y. Inverse design of materials by machine learning. *Materials*, 15(5), 2022. ISSN 1996-1944. doi: 10.3390/ma15051811. URL <https://www.mdpi.com/1996-1944/15/5/1811>.
- Wines, D., Xie, T., and Choudhary, K. Inverse design of next-generation superconductors using data-driven deep generative models, 2023.
- Xie, T. and Grossman, J. C. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Phys. Rev. Lett.*, 120: 145301, Apr 2018. doi: 10.1103/PhysRevLett.120.145301. URL <https://link.aps.org/doi/10.1103/PhysRevLett.120.145301>.
- Xie, T., Fu, X., Ganea, O.-E., Barzilay, R., and Jaakkola, T. S. Crystal diffusion variational autoencoder for periodic material generation. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=03RLpj-tc_.
- Xu, K., Zhang, M., Li, J., Du, S. S., Kawarabayashi, K.-I., and Jegelka, S. How neural networks extrapolate: From feedforward to graph neural networks. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=UH-cmocLJC>.
- Zhang, Y., Apley, D. W., and Chen, W. Bayesian optimization for materials design with mixed quantitative and qualitative variables. *Scientific Reports*, 10 (1):4924, Mar 2020. ISSN 2045-2322. doi: 10.1038/s41598-020-60652-9. URL <https://doi.org/10.1038/s41598-020-60652-9>.
- Zhao, Y., Al-Fahdi, M., Hu, M., Siriwardane, E. M., Song, Y., Nasiri, A., and Hu, J. High-throughput discovery of novel cubic crystal materials using deep generative neural networks. *Advanced Science*, 8(20):2100566, 2021.
- Zhao, Y., Siriwardane, E. M. D., Wu, Z., Fu, N., Al-Fahdi, M., Hu, M., and Hu, J. Physics guided deep learning for generative design of crystal materials with symmetry constraints. *npj Computational Materials*, 9 (1):38, Mar 2023. ISSN 2057-3960. doi: 10.1038/s41524-023-00987-9. URL <https://doi.org/10.1038/s41524-023-00987-9>.
- Zhu, J., Shen, Y., Zhao, D., and Zhou, B. In-domain GAN inversion for real image editing. In *Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVII*, pp. 592–608, Berlin, Heidelberg, 2020. Springer-Verlag. ISBN 978-3-030-58519-8. doi: 10.1007/978-3-030-58520-4_35. URL https://doi.org/10.1007/978-3-030-58520-4_35.

Zuo, Y., Qin, M., Chen, C., Ye, W., Li, X., Luo, J., and Ong, S. P. Accelerating materials discovery with Bayesian optimization and graph deep learning. *Materials Today*, 51:126–135, 2021. ISSN 1369-7021. doi: <https://doi.org/10.1016/j.mattod.2021.08.012>. URL <https://www.sciencedirect.com/science/article/pii/S1369702121002984>.

A. Supplemental figures




Figure 6. A PGCGM-generated structure for LiZn_2Pt with space group 225: $Fm\bar{3}m$ that is predicted by ALIGNN to have a negative decomposition enthalpy ΔH_d .

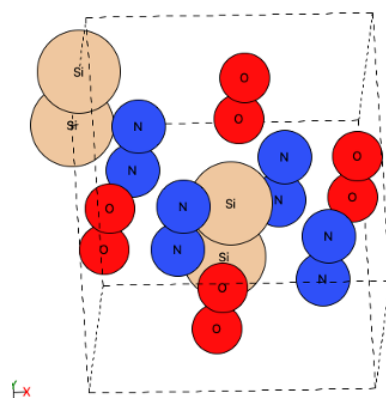


Figure 7. A PGCGM-generated structure for SiN_2O_2 with space group 62: $Pnma$ predicted by ALIGNN to have a positive decomposition enthalpy ΔH_d .

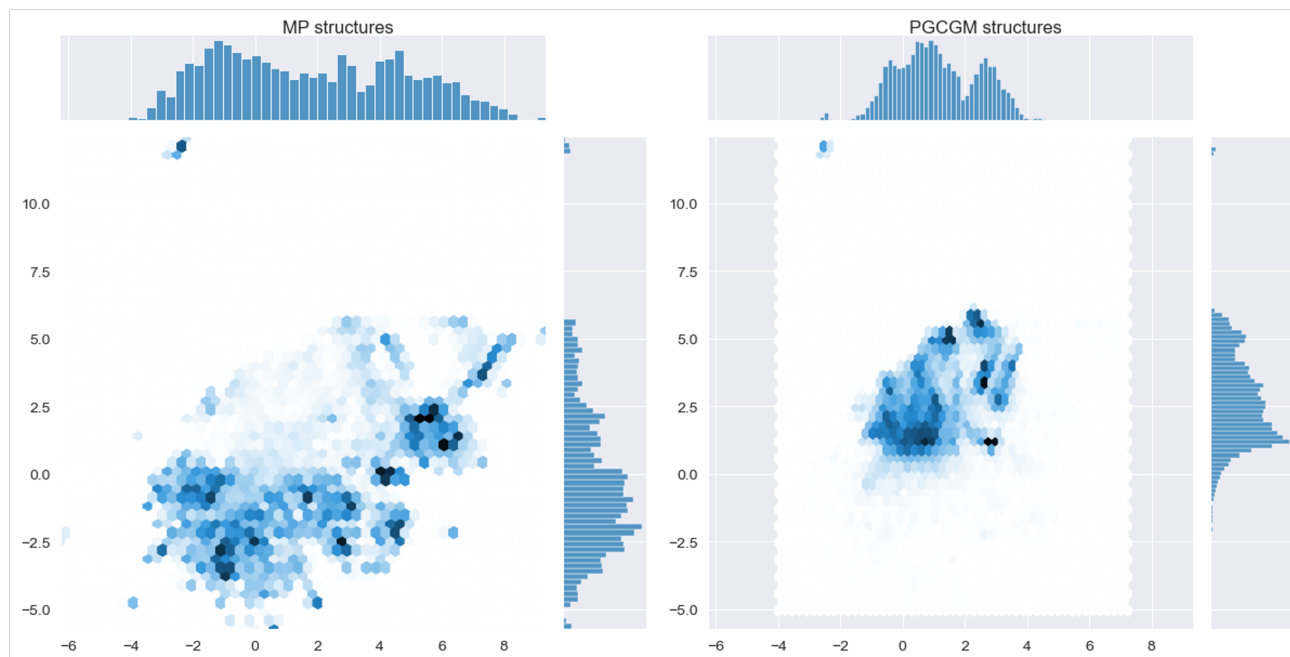


Figure 8. We use UMAP (McInnes et al., 2020) and project the ALIGNN embeddings of the 27,116 PGCGM (Zhao et al., 2023)-generated structures and the MP structures of the ALIGNN training data into two dimensions (shading indicates regions of high concentration). The PGCGM-generated structures are much more compressed in the latent space than the embeddings of the ALIGNN training data.




Figure 9. We repeat the distance-to-training-set analysis of Figure 4 while only using ALIGNN's (Choudhary & DeCost, 2021) training data. We hold out 20% (7,568 structures) of its training dataset, and we calculate the minimal distance from each point to the remaining 30,274 training points. Unlike Figure 4, there is no clear correlation between the ALIGNN's predicted decomposition enthalpy and the distance to the remaining training points. This suggests that the correlation between predicted decomposition enthalpy and distance-to-training set occurs specifically in the case of out-of-domain data.

B. Data details

For PGCGM (Zhao et al., 2023), we used the pre-trained model, which was trained on 33,172 ternary structures sourced from MP (Jain et al., 2013), ICSD (Belsky et al., 2002), and OQMD (Saal et al., 2013). The IDs were made available by the PGCGM authors on GitHub³. We trained our ALIGNN (Choudhary & DeCost, 2021) model on the 37,842 ternary structures from MP, selected from the 85,014 structures analyzed by Bartel et al. 2020.

³https://github.com/MilesZhao/PGCGM/blob/main/data/ids_for_mp_oqmd_icسد.csv