CompGen: A Conditional Generation Framework for Inverse Composition Design of Catalytic Surfaces

Anonymous Author(s)

Affiliation Address email

Abstract

Generating adsorption configurations, that is, how small atoms or molecules bind to complex catalyst surfaces, remains underexplored in inverse materials design. We present CompGen, a conditional generative framework that reformulates 3D structure prediction as a 2D shell-wise composition task centered on the adsorption site. CompGen uses a Chemically Informed Autoencoder (CIAE) to embed sparse compositions into a continuous, periodic table aware latent space learned with a multi-stage pretraining process. A conditional diffusion model then samples in this latent space under multi-physical conditions, including adsorbate identity, adsorption energy, and relevant elements, enabling inverse composition design of catalytic surfaces. Pretrained on a subset of Open Catalyst 2020, CompGen is fine-tuned to more complex high-entropy alloy (HEA) surfaces and achieves strong fine-tuned performance. Extensive experiments show robust zero-shot and few-shot behavior, highlighting CompGen's effectiveness for data-efficient, domain-transferable inverse design of catalytic surfaces.

1 Introduction

2

3

5

6

7

10

11

12

13

14

16

17

19

20

21

22

23

24

25

27

28

29

30

31

34

35

Inverse design of catalytic materials is a central AI-for-Science challenge with direct impact on energy and environmental technologies (Freeze et al., 2019; Noh et al., 2020; Wang et al., 2023). Catalytic performance is governed by active sites, the localized regions on a surface where reactions occur (Vogt and Weckhuysen, 2022). Despite rapid advances in generative models for inverse materials discovery (Gebauer et al., 2022; Xiao et al., 2023; Zheng et al., 2024; Park et al., 2024; Zeni et al., 2025; Joshi et al., 2025), applying them to catalytic surfaces remains difficult because local environments around active sites are uncertain, heterogeneous, and should be designed under constraints from key reactivity descriptors such as adsorption energy (Wang et al., 2025b). Existing approaches that generate explicit 3D coordinates also struggle to encode translational, rotational, and periodic symmetries efficiently (Kolluru et al., 2022; Duval et al., 2023; Wang and You, 2025). Recent equivariant diffusion methods, Reinforcement Learning schemes, and structure-search workflows have advanced adsorbate placement prediction and surface generation, but they mainly target simpler systems and often lack explicit, property-conditioned local microstructural design (Lacombe et al., 2023; Cornet et al., 2024; Kolluru and Kitchin, 2024; Rønne et al., 2024; Wang et al., 2025a).

In this work, we reformulate inverse design of catalytic surfaces as a simpler, yet chemically meaningful **composition generation task**. We focus on generating the proportions of each type of elements in two concentric spherical shells around a surface active site: *an inner shell 1* that interacts directly with the adsorbed atoms or molecules (i.e., the adsorbates) and *an outer shell 2* that surrounds the inner shell. This two-shell approach bypasses explicit 3D geometry, but preserves essential ensemble and ligand effects used to tune catalytic properties (Li et al., 2018; Pedersen et al., 2022). *shell 1* controls ensemble feasibility by fixing element counts near the adsorbate, while *shell 2* enables ligand tuning

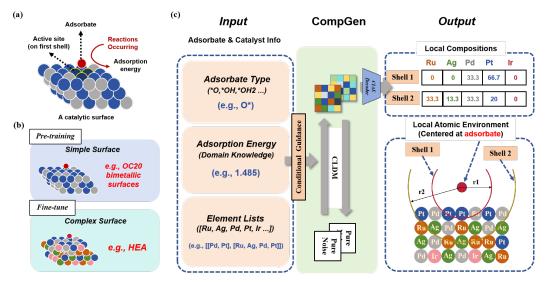


Figure 1: The illustration of (a) An example of adsorption configuration on a catalytic surface. (b) CompGen domain adaptation. (c) CompGen Pipeline.

through second shell composition. This setup provides useful priors that shrink the design space and reduce the number of candidate microstates, enabling more efficient high-throughput screening.

To address the composition generation task, we introduce **CompGen**, a conditional generative frame-39 40 work for inverse composition design task. CompGen couples a Chemically Informed Autoencoder 41 (CIAE) with a conditional diffusion model. CIAE projects shell-wise composition vectors into a chemically-aware latent space using a periodic table representation (PTR) prior (Feng et al., 2021) by 42 leveraging on a two-stage pretraining scheme. The conditional diffusion model then samples from 43 latent space under physical conditions, including categorical labels and numerical targets such as 44 adsorption energy. CompGen is pretrained on a subset of the Open Catalyst 2020 (OC20) dataset 45 (Chanussot et al., 2021), covering surfaces with up to three elements, and fine-tuned on a five-element 46 HEA dataset (Clausen et al., 2024), demonstrating strong zero-shot and few-shot generalization. 47 To the best of our knowledge, this is the first framework explicitly focused on inverse *composition* 48 design of catalytic surfaces. Our contributions are: (i) we cast 3D inverse surface design as a tractable 49 two-shell composition problem that naturally preserves ensemble and ligand effects; (ii) we develop 50 the CompGen framework, combining a periodic table-aware CIAE with conditional diffusion for 51 property-conditioned generation; and (iii) we demonstrate CompGen can robustly transfer from 52 simple surfaces to compositionally complex HEAs via parameter-efficient fine-tuning, and validate 53 the model architecture with comprehensive experiments. 54

2 Proposed CompGen Framework

2.1 Problem formulation

55

56

We cast inverse catalytic surfaces design task as generating two shell-wise compositions under given conditions. We define the local neighborhood of an adsorption site by two concentric shells of surface atoms around the adsorbate's *central* (binding) atom¹. (1) **First shell** C_1 : surface atoms within 2.5 Å cutoff radius of the central atom, i.e., its nearest neighbors by distance. (2) **Second shell** C_2 : surface atoms within 5.0 Å of the central atom that are the nearest neighbors of the first shell atoms.

For each shell $k \in \{1,2\}$, we use a composition vector $\mathbf{c}^k \in \mathbb{R}^D$ over D=118 chemical elements from the periodic table ordered by atomic number. The i-th entry $c_i^k \in [0,1]$ denotes the normalized atomic percentage of the i-th element within shell k, with $\sum_{i=1}^D c_i^k = 1$. We then stack the two shell compositions as $\mathbf{C} = [\mathbf{c}^1, \mathbf{c}^2] \in \mathbb{R}^{2 \times D}$ to be the generation target of CompGen framework.

¹The *central atom of the adsorbate* is the atom that forms the primary bond to the surface, e.g., 0 in *OH, C in *CH3, and N in *NH3.

We condition generation on three practical descriptors of an arbitrary pair of adsorbate-surface configuration: the adsorbate S, the adsorption energy E, and an element list L that specifies the allowed element types in each shell. In our simplified setting, these three descriptors compactly define the adsorption model as presented in Figure 1.

Taken together, formally, we aim to learn the conditional distribution of the surface composition $p(C_1, C_2|S, E, L)$, so that we can sample novel, constraint-consistent shell compositions. Our proposed CompGen framework reduces the search complexity by at least a polynomial factor in the number of sites. See more details in Appendix D.

2.2 Chemically Informed Autoencoder (CIAE)

74

98

We design the CIAE by imposing chemical prior information in two ways: (i) it represents each shell composition on a periodic table grid inspired by prior works (Zheng et al., 2018; Feng et al., 2021), and (ii) it uses a two-stage pretraining scheme that first passes property-aware features into the encoder and then trains the full autoencoder for reconstruction.

To obtain a compact latent space that preserves core chemical regularities, we pre-process each shell composition $\mathbf{c}^k \in \mathbb{R}^D$ by mapping it to a 2D grid representation \mathbf{P} that encodes the periodic table's layout. We then stack the two shell PTRs to form a two-channel input tensor $\mathbf{x}_0 = [\mathbf{P}_1; \mathbf{P}_2] \in \mathbb{R}^{2 \times H \times W}$. See more details in Appendix \mathbf{B} .

Stage I: Supervised Encoder Pretraining. To inject property awareness into the latent space, we adopt the General and Transferable Deep Learning (GTDL) framework (Feng et al., 2021) by using a VGG-like CNN backbone pretrained on the Materials Project dataset (Jain et al., 2013) for a binary classification task to serve as our encoder $f_{\rm enc}$. It maps $\mathbf{x_0}$ to a latent vector $\mathbf{z} = f_{\rm enc}(\mathbf{x_0}; \theta_{\rm enc}) \in \mathbb{R}^{2 \times d_z \times d_z}$. However, unlike the original framework, we do not augment PTRs with handcrafted features, as our objective is only for composition reconstruction.

Stage II: Pretraining the Autoencoder. With the stage I encoder as initialization, we train the full autoencoder end-to-end for high-fidelity reconstruction of shell-wise compositions. As mentioned in stage I, the stacked PTR tensor $\mathbf{x_0} \in \mathbb{R}^{2 \times H \times W}$ is encoded to a low-dimensional latent vector \mathbf{z} . In stage II, a decoder f_{dec} maps \mathbf{z} back to the composition space, producing $\hat{\mathbf{C}} = f_{\text{dec}}(\mathbf{z}; \theta_{\text{dec}}) \in \mathbb{R}^{2 \times D}$, i.e., per-shell elemental fractions rather than PTR pixels. The decoder uses a convolutional block with a residual connection followed by fully connected layers. The training objective is to minimize the reconstruction error between the original composition tensor \mathbf{C} and the reconstructed tensor $\hat{\mathbf{C}}$. We optimize $(\theta_{\text{enc}}, \theta_{\text{dec}})$ with a *Mean Squared Error* (MSE) loss between the ground-truth compositions \mathbf{C} and reconstructions $\hat{\mathbf{C}}$, averaged over shells and elements.

2.3 Conditional Latent Diffusion Model

The chemically informed latent space learned by CIAE in pretraining stages provides the foundation 99 for the generation task. Given a latent embedding z from the CIAE encoder, we train a conditional 100 diffusion model to learn the distribution of z conditioned on: (1) Adsorbate S. We encode S as a 101 one-hot vector of dimension K_s ($K_s = 13$ in our experiments, covering 13 types of the most basic 102 adsorbates such as *0, *CH, *NH2) from the OC20 dataset. A learned embedding layer maps this to 103 $\mathbf{s} \in \mathbb{R}^{d_h}$. (2) Adsorption energy E. The scalar target energy (in unit of eV) is repeated to length K_s for the stability of training and then passed through a MLP to produce $e \in \mathbb{R}^{d_h}$. (3) Element list L. 105 The allowed element types (e.g., [Ru, Pt, Pd, Ag, Ir]) of each shell are embedded to yield $l \in \mathbb{R}^{d_h}$. 106 **Latent Diffusion.** Let $z_0 = z$ denote the latent representation produced by the pretrained CIAE, 107 we implement the standard CLDM process. The forward process corrupts latent z_0 into z_t over 108 discrete timesteps $t \in \{0,...,T\}$ by gradually adding Gaussian noise. T is chosen large enough 109 that z_T is approximately standard normal. A neural network ϵ_{θ} is trained to run the reverse process, 110 denoising from t = T to t = 0 by predicting the additive noise. The training objective minimizes the expected MSE between the true noise and the network's prediction over noisy latents \mathbf{z}_t : $\mathcal{L}(\theta)$ 112 $\mathbb{E}_{\mathbf{z}_0, \epsilon \sim \mathcal{N}(0, I), t} \left[\|\epsilon - \epsilon_{\theta} \left(\mathbf{z}_t, t, \mathbf{s}, \mathbf{e}, \mathbf{l} \right) \|^2 \right]$, where ϵ_{θ} is conditioned on $(\mathbf{s}, \mathbf{e}, \mathbf{l})$.

3 Experiment

Metrics. We evaluate the performance of CompGen with three complementary metrics: (i) *Fréchet Distance* (FD) for distributional similarity of generated CIAE latent embedding, (ii) *Leakage* for

Table 1: Performance comparison of model architecture and fine-tuning methods.

more in remainmente comput						
Experiments on Shell 1	OC20			HEA		
	FD↓	Leakage↓	Sim ↑	FD↓	Leakage↓	Sim ↑
DiT Baseline Pretrain	137.22	20.44	0.11	143.46	20.46	0.10
DiT Baseline + LoRA Fine-tune	_	-	-	121.52	20.39	0.11
U-Net Pretrain	0.48	0.03	0.92	0.43	0.05	0.93
U-Net + Full-param Fine-tune	_	_	-	0.23	0.03	0.95
U-Net + LoRA Fine-tune	_	-	-	0.27	0.04	0.94
DiT Pretrain	0.71	0.13	0.90	2.55	0.18	0.90
DiT + Full-param Fine-tune	_	_	-	0.25	0.04	0.95
DiT + LoRA Fine-tune	_	_	_	0.86	0.12	0.93
E-marinanta an Chall 2		OC20			HEA	
Experiments on Shell 2	FD↓	OC20 Leakage↓	Sim ↑	FD↓	HEA Leakage↓	Sim ↑
Experiments on Shell 2 DiT Baseline Pretrain	FD↓ 26.24		Sim ↑ -0.11	FD↓		Sim ↑ -0.13
		Leakage↓	'		Leakage↓	· ·
DiT Baseline Pretrain		Leakage↓	'	43.48	Leakage↓ -11.13	-0.13
DiT Baseline Pretrain DiT Baseline + LoRA Fine-tune	26.24	Leakage↓ -11.30 -	-0.11	43.48 46.14	Leakage↓ -11.13 -11.31	-0.13 -0.15
DiT Baseline Pretrain DiT Baseline + LoRA Fine-tune U-Net Pretrain	26.24	Leakage↓ -11.30 -	-0.11	43.48 46.14 11.72	Leakage↓ -11.13 -11.31 0.34	-0.13 -0.15 0.48
DiT Baseline Pretrain DiT Baseline + LoRA Fine-tune U-Net Pretrain U-Net + Full-param Fine-tune	26.24	Leakage↓ -11.30 -	-0.11	43.48 46.14 11.72 2.41	Leakage↓ -11.13 -11.31 0.34 0.31	-0.13 -0.15 0.48 0.52
DiT Baseline Pretrain DiT Baseline + LoRA Fine-tune U-Net Pretrain U-Net + Full-param Fine-tune U-Net + LoRA Fine-tune	26.24 - 1.79 - -	Leakage↓ -11.30 - 0.21	-0.11 - 0.71 - -	43.48 46.14 11.72 2.41 2.80	Leakage↓ -11.13 -11.31 0.34 0.31 0.31	-0.13 -0.15 0.48 0.52 0.59

compliance with compositional constraints, and (iii) *Cosine Similarity* for point-wise measure of quality for the final generated two-shell composition (see Appendix G.2).

Baseline. To form a chemically agnostic baseline, the composition vector $\mathbf{c} \in \mathbb{R}^{118}$ is reshaped into a spatial grid and zero-padded to create a tensor $X \in \mathbb{R}^{48 \times 48}$. This tensor is fed directly to the diffusion model, bypassing the CIAE entirely (see Appendix G.1).

CIAE Pretraining. We pretrain the CIAE on 95% of the selected OC20 subset to learn a chemically structured latent space from shell-wise compositions. We report (i) MSE of latent reconstruction and (ii) point-wise Cosine Similarity between reconstructed and ground truth 118-dimensional composition vectors for shell 1 and shell 2. Evaluation is performed on held-out OC20 and out-of-domain HEA test sets. Results are presented in Table 3 (Appendix G.3).

OC20 Pretraining and Fine-tune on HEA. We pretrain CompGen on the OC20 subset to evaluate both U-Net and DiT backbones with the given three conditions (S, E, and L) in the chemically informed latent space. To assess the transferability and adaptability of CompGen, we fine-tune the model on the HEA dataset whose compositional and structural characteristics differ markedly from OC20. We consider two strategies: (i) *Low-Rank Adaptation* (LoRA) (Hu et al., 2022) and (ii) *full-parameter fine-tuning*.

Results and Analysis. Table 1 reveals four conclusions. (i) A chemistry-aware latent space is crucial: the chemically agnostic baseline fails to learn a meaningful distribution, whereas our pretrained models achieve better performance (FD). (ii) Both backbones exhibit notable zero-shot transferability to HEA; the U-Net even improves on shell 1 metrics without fine-tuning. (iii) Fine-tuning consistently boosts performance as evaluated by FD, with full-parameter fine-tuning of the U-Net delivering the best results across both shells. (iv) Parameter-efficient LoRA is a strong alternative that can match or exceed the performance of full-parameter fine-tuning. Especially on the more challenging shell 2 composition generation task, it achieves the highest Sim score.

4 Conclusion

In this work, we cast inverse catalysts design as a two-shell composition generation task and introduce CompGen, a modular framework that couples CIAE with CLDM to supply compositionally actionable priors that collapse the otherwise vast 3D search space. Across multiple metrics, CompGen achieves high generation quality and strong transfer from subset of OC20 surfaces to compositionally complex HEA surfaces via efficient fine-tuning. The framework is architecture-agnostic and can accommodate alternative encoders, generators, and conditioning schemes as the task evolves.

References

- AI4Science, M., Hernandez-Garcia, A., Duval, A., Volokhova, A., Bengio, Y., Sharma, D., Carrier,
 P. L., Benabed, Y., Koziarski, M., and Schmidt, V. (2023). Crystal-gfn: sampling crystals with
 desirable properties and constraints. *arXiv preprint arXiv:2310.04925*.
- 152 Chanussot, L., Das, A., Goyal, S., Lavril, T., Shuaibi, M., Riviere, M., Tran, K., Heras-Domingo, J., 153 Ho, C., Hu, W., et al. (2021). Open catalyst 2020 (oc20) dataset and community challenges. *Acs* 154 *Catalysis*, 11(10):6059–6072.
- 155 Clausen, C. M., Rossmeisl, J., and Ulissi, Z. W. (2024). Adapting oc20-trained equiformerv2 models 156 for high-entropy materials. *The Journal of Physical Chemistry C*, 128(27):11190–11195.
- Cornet, F., Benediktsson, B., Hastrup, B., Schmidt, M. N., and Bhowmik, A. (2024). Om-diff:
 inverse-design of organometallic catalysts with guided equivariant denoising diffusion. *Digital Discovery*, 3(9):1793–1811.
- Duval, A., Mathis, S. V., Joshi, C. K., Schmidt, V., Miret, S., Malliaros, F. D., Cohen, T., Lio, P.,
 Bengio, Y., and Bronstein, M. (2023). A hitchhiker's guide to geometric gnns for 3d atomic
 systems. arXiv preprint arXiv:2312.07511.
- Feng, S., Fu, H., Zhou, H., Wu, Y., Lu, Z., and Dong, H. (2021). A general and transferable deep
 learning framework for predicting phase formation in materials. *npj Computational Materials*,
 7(1):10.
- Freeze, J. G., Kelly, H. R., and Batista, V. S. (2019). Search for catalysts by inverse design: artificial intelligence, mountain climbers, and alchemists. *Chemical reviews*, 119(11):6595–6612.
- Gebauer, N. W., Gastegger, M., Hessmann, S. S., Müller, K.-R., and Schütt, K. T. (2022). Inverse design of 3d molecular structures with conditional generative neural networks. *Nature communications*, 13(1):973.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. (2017). Gans trained by a
 two time-scale update rule converge to a local nash equilibrium. *Advances in Neural Information Processing Systems*, 30.
- Hoogeboom, E., Satorras, V. G., Vignac, C., and Welling, M. (2022). Equivariant diffusion for
 molecule generation in 3d. In *International conference on machine learning*, pages 8867–8887.
 PMLR.
- Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., et al. (2022). Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.
- Jain, A., Ong, S. P., Hautier, G., Chen, W., Richards, W. D., Dacek, S., Cholia, S., Gunter, D., Skinner, D., Ceder, G., et al. (2013). Commentary: The materials project: A materials genome approach to accelerating materials innovation. *APL materials*, 1(1).
- Jiao, R., Huang, W., Lin, P., Han, J., Chen, P., Lu, Y., and Liu, Y. (2023). Crystal structure prediction by joint equivariant diffusion. *Advances in Neural Information Processing Systems*, 36:17464–17497.
- Joshi, C. K., Fu, X., Liao, Y.-L., Gharakhanyan, V., Miller, B. K., Sriram, A., and Ulissi, Z. W. (2025).
 All-atom diffusion transformers: Unified generative modelling of molecules and materials. *arXiv* preprint arXiv:2503.03965.
- Kempen, L. H. and Andersen, M. (2025). Inverse catalysts: tuning the composition and structure of oxide clusters through the metal support. *npj Computational Materials*, 11(1):8.
- Kim, S., Noh, J., Gu, G. H., Aspuru-Guzik, A., and Jung, Y. (2020). Generative adversarial networks for crystal structure prediction. *ACS central science*, 6(8):1412–1420.
- Kohn, W. and Sham, L. J. (1965). Self-consistent equations including exchange and correlation effects. *Physical review*, 140(4A):A1133.

- Kolluru, A. and Kitchin, J. R. (2024). Adsorbdiff: Adsorbate placement via conditional denoising diffusion. *arXiv preprint arXiv:2405.03962*.
- Kolluru, A., Shuaibi, M., Palizhati, A., Shoghi, N., Das, A., Wood, B., Zitnick, C. L., Kitchin, J. R.,
 and Ulissi, Z. W. (2022). Open challenges in developing generalizable large-scale machine-learning
 models for catalyst discovery. ACS Catalysis, 12(14):8572–8581.
- Lacombe, R., Hendren, L., and El-Awady, K. (2023). Adsorbrl: Deep multi-objective reinforcement learning for inverse catalysts design. *arXiv* preprint arXiv:2312.02308.
- Li, H., Shin, K., and Henkelman, G. (2018). Effects of ensembles, ligand, and strain on adsorbate binding to alloy surfaces. *The Journal of chemical physics*, 149(17).
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. (2015). Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533.
- Mok, D. H. and Back, S. (2024). Generative pretrained transformer for heterogeneous catalysts. *Journal of the American Chemical Society*, 146(49):33712–33722.
- Noh, J., Gu, G. H., Kim, S., and Jung, Y. (2020). Machine-enabled inverse design of inorganic solid materials: promises and challenges. *Chemical Science*, 11(19):4871–4881.
- Nouira, A., Sokolovska, N., and Crivello, J.-C. (2018). Crystalgan: learning to discover crystallographic structures with generative adversarial networks. *arXiv preprint arXiv:1810.11203*.
- Pakornchote, T., Choomphon-Anomakhun, N., Arrerut, S., Atthapak, C., Khamkaeo, S., Chotibut, T., and Bovornratanaraks, T. (2024). Diffusion probabilistic models enhance variational autoencoder for crystal structure generative modeling. *Scientific Reports*, 14(1):1275.
- Park, H., Li, Z., and Walsh, A. (2024). Has generative artificial intelligence solved inverse materials design? *Matter*, 7(7):2355–2367.
- Pedersen, J. K., Clausen, C. M., Skjegstad, L. E. J., and Rossmeisl, J. (2022). A mean-field model for oxygen reduction electrocatalytic activity on high-entropy alloys. *ChemCatChem*, 14(18):e202200699.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. (2019). Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Rønne, N., Aspuru-Guzik, A., and Hammer, B. (2024). Generative diffusion model for surface structure discovery. *Physical Review B*, 110(23):235427.
- Vogt, C. and Weckhuysen, B. M. (2022). The concept of active site in heterogeneous catalysis. *Nature Reviews Chemistry*, 6(2):89–111.
- Wang, B., Zheng, S., Wu, J., Li, J., and Pan, F. (2025a). Inverse design of catalytic active sites via interpretable topology-based deep generative models. *npj Computational Materials*, 11(1):147.
- Wang, H., Fu, T., Du, Y., Gao, W., Huang, K., Liu, Z., Chandak, P., Liu, S., Van Katwyk, P., Deac, A., et al. (2023). Scientific discovery in the age of artificial intelligence. *Nature*, 620(7972):47–60.
- Wang, Z., Li, W., Wang, S., and Wang, X. (2025b). The future of catalysis: Applying graph neural
 networks for intelligent catalyst design. Wiley Interdisciplinary Reviews: Computational Molecular
 Science, 15(2):e70010.
- Wang, Z. and You, F. (2025). Leveraging generative models with periodicity-aware, invertible and
 invariant representations for crystalline materials design. *Nature Computational Science*, pages
 1–12.
- Xiao, H., Li, R., Shi, X., Chen, Y., Zhu, L., Chen, X., and Wang, L. (2023). An invertible, invariant
 crystal representation for inverse design of solid-state materials using generative deep learning.
 Nature Communications, 14(1):7027.

- Zeni, C., Pinsler, R., Zügner, D., Fowler, A., Horton, M., Fu, X., Wang, Z., Shysheya, A., Crabbé, J.,
 Ueda, S., et al. (2025). A generative model for inorganic materials design. *Nature*, 639(8055):624–632.
- Zhao, Y., Al-Fahdi, M., Hu, M., Siriwardane, E. M., Song, Y., Nasiri, A., and Hu, J. (2021). High throughput discovery of novel cubic crystal materials using deep generative neural networks.
 Advanced Science, 8(20):2100566.
- Zheng, S., He, J., Liu, C., Shi, Y., Lu, Z., Feng, W., Ju, F., Wang, J., Zhu, J., Min, Y., et al. (2024).
 Predicting equilibrium distributions for molecular systems with deep learning. *Nature Machine Intelligence*, 6(5):558–567.
- Zheng, X., Zheng, P., and Zhang, R.-Z. (2018). Machine learning material properties from the periodic table using convolutional neural networks. *npj Computational Materials*, 4(1):1–7.

A Related work

Inverse materials design. Generative models have advanced inverse materials design by mapping target properties to novel structures. Early GAN-based methods (Nouira et al., 2018; Kim et al., 2020; Zhao et al., 2021) laid the groundwork, followed by diffusion-based (Hoogeboom et al., 2022; Jiao et al., 2023; Pakornchote et al., 2024) and flow-based approaches (AI4Science et al., 2023) that further improved the task with high-fidelity generalization. Recent representative diffusion models include MatterGen (Zeni et al., 2025), which uses property-conditioned denoising to improve the generation of stable, unique, and novel inorganic materials, and All-atom Diffusion Transformer (ADiT) (Joshi et al., 2025), which maps molecules and crystals into a shared VAE latent space and jointly generates periodic and non-periodic atomic structures via a latent-space diffusion Transformer.

Inverse catalysts design. Inverse design of catalytic surfaces raises distinct challenges beyond bulk generation. A range of paradigms has been explored: reinforcement learning, language-model-based generation, and diffusion models. AdsorbRL (Lacombe et al., 2023) employs a Deep Q-Network (Mnih et al., 2015) to navigate vast compositional spaces and identify catalysts optimized with given adsorption energy targets. CatGPT (Mok and Back, 2024) leverages a GPT-2-based language model (Radford et al., 2019) to generate string-based representations of catalyst surfaces, enabling fine-tuning for downstream property prediction tasks validated by Density Functional Theory (DFT) (Kohn and Sham, 1965) calculations. Among diffusion models, Rønne et al. (Rønne et al., 2024) propose a rotationally equivariant diffusion framework with force-field guidance to sample low-energy silver-oxide surface structures, though without explicit conditioning on target properties. OM-Diff (Cornet et al., 2024) introduces a guided equivariant diffusion model that generates 3D structures of homogeneous organometallic catalysts, conditioned on specific metal centers via regressor-informed denoising. AdsorbDiff (Kolluru and Kitchin, 2024) predicts the optimal adsorbate binding orientations and placements on catalytic surfaces. However, it focuses on improving placement success rates rather than generating the local surface environment of active sites. Complementary to these generative efforts, structure-search workflows have been proposed for inverse catalysts such as metal-oxide interfaces (Kempen and Andersen, 2025). An thermodynamics-guided search over Zn_yO_x and In_yO_x clusters on pure metal surfaces uncovers stable active site motifs and highlights the importance of site diversity in design. Most recently, PGH-VAE (Wang et al., 2025a) applies a topology-aware VAE with features from topological algebraic analysis to inverse design of active sites on IrPdPtRhRu HEA surfaces.

B Periodic table representation (PTR) details

Our approach begins with how we represent the atomic environment in the model. Instead of a simple list of atoms, we leverage the inherent structure of the periodic table to create a chemically-aware input representation. The compositional vector for each shell mentioned in Probelm Setup, $\mathbf{c}^k \in \mathbb{R}^D$, is transformed into a 2D grid-like representation. This is achieved by mapping the fractional concentration of each element to a specific location on a 2D grid that mirrors the layout of the periodic table. Formally, the entry at position (h,w) for shell k is defined as: $(\mathbf{P}_k)_{h,w} = c_i^k$ if cell (h,w) corresponds to element i and 0 otherwise, where \mathbf{P}_k is called the PTR for shell k. Cells for absent elements and empty positions are zero-padded. Specifically, we construct $\mathbf{P}_k \in \mathbb{R}^{H \times W}$, with dimensions corresponding to the periodic table (i.e., H=9,W=18). Each cell (h,w) in this grid is uniquely assigned to a chemical element based on its position. The value of each cell is then fitted with the fractional concentration of its corresponding element from the composition vector \mathbf{c}^k . Formally, the entry at position (h,w) for shell k is defined as:

$$(\mathbf{P}_k)_{h,w} = \begin{cases} c_i^k & \text{if cell } (h,w) \text{ corresponds to element } i \\ 0 & \text{otherwise} \end{cases}$$

where \mathbf{P}_k is called the PTR for shell k. This representation transforms the compositional data into a grid-like tensor that spatially encodes chemical relationships. For instance, elements in the same group appear in the same column, while elements in the same period appear in the same row. This structure provides a powerful inductive bias, rendering a 2D Convolutional Neural Network (CNN) well-suited architecture for our autoencoder. This allows the model to recognize patterns of chemically similar elements, effectively capturing complex relationships that are missing with non-spatial representation.

C Analysis of CompGen data efficiency for training

To evaluate the data efficiency of our model, we conduct a systematic study using 40%, 60%, and 80% of the training data. These results are compared to a backbone model trained on the full dataset (100%). The model performance is evaluated on both OC20 and HEA tasks, by using MSE between latent embeddings and cosine similarity between shell compositions. Results are shown in Table 2 and figure 2.

Table 2: Performance under varying proportion of training data. Lower MSE and higher similarity values indicate better performance.

Data	MSE (OC20) ↓	OC20 shell 1 ↑	OC20 shell 2 ↑	MSE (HEA) ↓	HEA shell 1 ↑	HEA shell 2↑
40%	31.7891	0.8640	0.8207	48.7329	0.8484	0.5634
60%	28.9034	0.8848	0.8288	46.1000	0.8662	0.5715
80%	29.5925	0.8813	0.8343	44.3454	0.9014	0.5776
100%	26.1264	0.9037	0.8107	27.8641	0.8988	0.5627

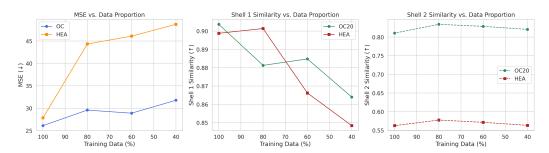


Figure 2: **Data scaling effect across multiple metrics.** Performance of the model under varying proportion of training data (40%, 60%, 80%, 100%). (**Left**) Mean squared error (MSE) for both OC20 and HEA tasks. (**Middle**) Similarity of shell 1 composition across generated structures. (**Right**) Similarity of shell 2 composition. Results demonstrate consistent model improvement with more training data, following typical deep learning scaling laws.

From Table 2, we observe a general trend between training data size and model performance. On the OC20 task, the MSE steadily improves from 31.79 at 40% to 28.90 at 60%, and reaches 29.59 at 80%, approaching the full-data baseline of 26.12. The OC20 shell similarity scores also improve, with OC20 shell 2 increasing from 0.8207 to 0.8343, demonstrating better local structural consistency.

Similarly, on the HEA task, the MSE drops from 48.73 (40%) to 44.34 (80%), and HEA shell 2 similarity grows from 0.5634 to 0.5776. Notably, this slightly exceeds the backbone score (0.5627), suggesting that fine-tuning on a subset can potentially outperform full-data pretraining under certain domain shifts.

Overall, these results are consistent with the expected data scaling behavior in deep models. Model performance improves sublinearly with more data, and even partial training (60 - 80%) is sufficient to recover most of the backbone performance. This demonstrates the data efficiency and generalizability of the proposed method.

D CompGen efficiency analysis

Setup and notation. Fix an adsorption site type (e.g., atop, bridge, hollow) on a specified facet. Let n_1 be the number of *labeled* first shell sites (directly coordinated surface atoms; e.g., $n_1 = 1, 2, 3$ for atop/bridge/hollow respectively), and let n_2 be the number of *labeled* second shell sites (near-surface/subsurface neighbors within a fixed cutoff). Let \mathcal{E} be the full element set under consideration (e.g., HEA elements), with $|\mathcal{E}| = N_E$. For each shell $s \in \{1,2\}$, let $\mathcal{A}_s \subseteq \mathcal{E}$ denote the *allowed element set* for that shell (provided by design rules or domain constraints), with size $|\mathcal{A}_s| = N_s$. A *microstate* is a distinct assignment of elements to the n_s labeled sites of shell s (and likewise for the other shell); we count microstates for the two shells jointly by multiplication.

We consider three levels of constraints:

- 1. Unconstrained (U): Any site may take any element in \mathcal{E} .
 - 2. Subset-constrained (S): Shell s may take only elements in A_s (no composition counts
 - 3. COMPGEN-constrained (C): In addition to A_s , shell s has a composition count vector $\mathbf{K}_s = (K_{s,i})_{i \in A_s}$ with $K_{s,i} \in \mathbb{Z}_{\geq 0}$ and $\sum_{i \in A_s} K_{s,i} = n_s$.
- Unless stated, we treat sites as labeled.² 334
- Case U: Unconstrained. Each of the n_1 first shell sites and n_2 second shell sites may independently 335 take any of the N_E elements: 336

$$\#\mathcal{M}_{U} = N_{E}^{n_1 + n_2}. \tag{1}$$

Case S: Subset-constrained by shell. Restricting to A_1 and A_2 yields

$$\#\mathcal{M}_{S} = N_{1}^{n_{1}} N_{2}^{n_{2}}, \tag{2}$$

- which is a factor of $\left(\frac{N_1}{N_P}\right)^{n_1} \left(\frac{N_2}{N_P}\right)^{n_2}$ reduction relative to (1). 338
- Case C: COMPGEN composition-constrained. Given count vectors K_1 and K_2 , the number of 339 assignments for each shell is a multinomial coefficient:

$$\#\mathcal{M}_{\mathbf{C}} = \underbrace{\frac{n_1!}{\prod_{i \in \mathcal{A}_1} K_{1,i}!}}_{\text{first shell}} \times \underbrace{\frac{n_2!}{\prod_{i \in \mathcal{A}_2} K_{2,i}!}}_{\text{second shell}}.$$
 (3)

- This follows from counting the permutations of n_s labeled sites subject to exact element counts $K_{s,i}$ in shell s. 342
- **Reduction factors.** Equations (1), (2), (3) imply the hierarchy 343

$$\#\mathcal{M}_{\mathrm{C}} \leq \#\mathcal{M}_{\mathrm{S}} \leq \#\mathcal{M}_{\mathrm{U}}$$
.

In particular, 344

329

330 331

332 333

$$\frac{\#\mathcal{M}_{\mathrm{C}}}{\#\mathcal{M}_{\mathrm{S}}} \; = \; \frac{n_{1}!}{\prod_{i \in \mathcal{A}_{1}} K_{1,i}!} \; \frac{n_{2}!}{\prod_{i \in \mathcal{A}_{2}} K_{2,i}!} \; \cdot \; \frac{1}{N_{1}^{n_{1}} N_{2}^{n_{2}}}, \qquad \frac{\#\mathcal{M}_{\mathrm{C}}}{\#\mathcal{M}_{\mathrm{U}}} \; = \; \frac{n_{1}!}{\prod_{i} K_{1,i}!} \; \frac{n_{2}!}{\prod_{i} K_{2,i}!} \; \cdot \; \frac{1}{N_{E}^{n_{1}+n_{2}}},$$

- showing that COMPGEN yields exponential-in- n_s reductions when n_s is moderate (with exact factors governed by the multinomial denominators).
- **Concrete HEA example with motif-aware** n_1 . Consider a quinary HEA with \mathcal{E} 347
- $\{Ag, Ir, Pd, Pt, Ru\}$ $(N_E = 5)$. Fix a hollow adsorption motif on an fcc(111) facet so $n_1 = 3$ 348 (three directly coordinated surface atoms).³ Let the near-surface cutoff yield $n_2 = 6$ second 349
- shell sites (adjustable to your geometry). Assume design constraints: $A_1 = \{Pt, Pd\}$ $(N_1=2)$,
- $A_2 = \{ \text{Pt}, \text{Pd}, \text{Ag}, \text{Ru} \} (N_2 = 4).$ 351

Unconstrained:
$$\#\mathcal{M}_{U} = 5^{n_1+n_2} = 5^9 = 1,953,125.$$

Subset-constrained:
$$\#M_S = 2^3 \cdot 4^6 = 8 \cdot 4,096 = 32,768$$
.

Suppose COMPGEN proposes the first shell counts $\mathbf{K}_1=(K_{1,\mathrm{Pt}},K_{1,\mathrm{Pd}})=(2,1)$ and the second shell counts $\mathbf{K}_2=(K_{2,\mathrm{Pt}},K_{2,\mathrm{Pd}},K_{2,\mathrm{Ag}},K_{2,\mathrm{Ru}})=(3,1,1,1)$. Then

$$\#\mathcal{M}_{C} = \underbrace{\frac{3!}{2! \, 1!}}_{-3} \times \underbrace{\frac{6!}{3! \, 1! \, 1! \, 1!}}_{-120} = 360.$$

- Thus, relative to the unconstrained case, the search is reduced by a factor of $1.953,125/360 \approx 5.425$, 354
- and relative to the subset-only case by $32,768/360 \approx 91$. These counts are before any geometric 355
- relaxation or symmetry pruning. 356

²Labeled sites reflect distinct geometric positions around the adsorption center (e.g., the three specific metal atoms forming an fcc/hcp hollow). If certain positions are symmetry-equivalent, one may divide by the symmetry group size to obtain a reduced count; our formulas give an upper bound that is sufficient to show the reduction

³For atop and bridge motifs, set $n_1=1$ or 2.

Remarks on symmetry and unlabeled variants. If certain labeled positions are symmetry-equivalent (e.g., the three hollow sites under a C_3 rotation), one may divide the counts by the appropriate group size to obtain a tighter estimate. Alternatively, if one prefers to treat sites as unlabeled, the first shell count reduces to the number of distinct compositions only (one per feasible \mathbf{K}_1), which is $\binom{n_1+N_1-1}{N_1-1}$; COMPGEN then selects a single \mathbf{K}_1 , and the count is 1 for that shell. Our labeled-site model provides a conservative (larger) count and hence a conservative reduction factor.

E Dataset details

Two datasets are considered during the pretraining and fine-tuning stages of CompGen. For pretraining, we utilize a subset of the per-adsorbate trajectories from the OC20 dataset⁴. Specifically, we extract trajectories involving 13 common adsorbates (*0, *H, *0H, *0H2, *C, *CH, *CH2, *CH3, *CH4, *N, *NH, *NH2, and *NH3)⁵, with well-defined central atoms as reference to calculate the distance to the surface atoms for each shell. These adsorbates are adsorbed on distinct surfaces, each composed of at most three different element types. In total, 52 unique surface elements are presented, and the final dataset consists of 131,740 adsorption configurations. We first randomly split the dataset with 5% as test set, and further split the remaining 95% of the dataset into training and validation set with ratio of 90:10.

HEAs have recently gained attention as promising catalysts due to their highly diverse local structural composition, which enable fine-tuning of catalytic properties across a broad design space. To evaluate CompGen's fine-tuning capabilities, we adopt the HEA dataset from Clausen et al. (Clausen et al., 2024), which includes 4,892 distinct adsorption configurations of *0H and *0 on HEA surfaces composed of five elements: Ag, Ir, Pd, Pt, and Ru. The same pre-processing protocol is applied as in the pretraining dataset by extracting normalized first and second shell element compositions. We follow the original data split as used in the original paper, keeping the same 80:10:10 ratio for the fine-tuning experiments of CompGen.

We further provide more details about the OC20 dataset used in CompGen pretraining stage. First we have the heatmap for composition distribution over element vs. adsorption energy on both shell over adsorbate *C. The composition heatmap for OC20 is presented in Figure 3 and 4.

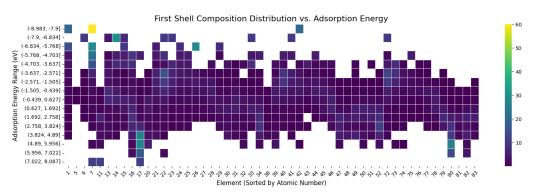


Figure 3: Composition heatmap for shell 1.

To further investigate the relationship between composition and adsorption energy, we perform a detailed analysis on a narrow slice of the property landscape. Figure 5 and Figure 6 specifically examine the "middle energy bin", visualizing the element compositions of the extracted 15 surfaces that exhibit the lowest adsorption energies within this range. This examination provides a high resolution of the diversity for element combinations that condition on a constrained range of target adsorption energies.

⁴The OC20 dataset is available at https://fair-chem.github.io/catalysts/datasets/oc20.html

⁵The * sign indicates that the atom or molecule is bound to the surface of the catalyst.

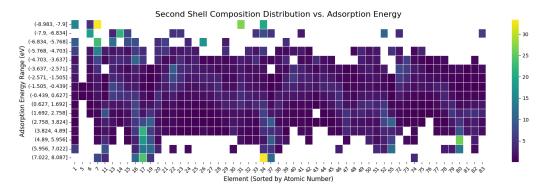


Figure 4: Composition heatmap for shell 2.

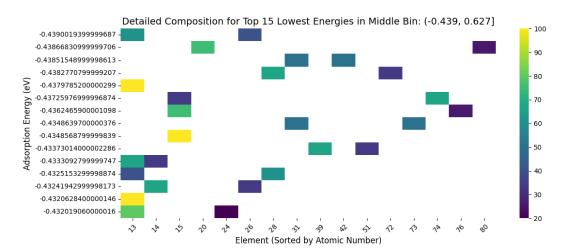


Figure 5: Composition heatmap for specific range of adsorption energies on shell 1.

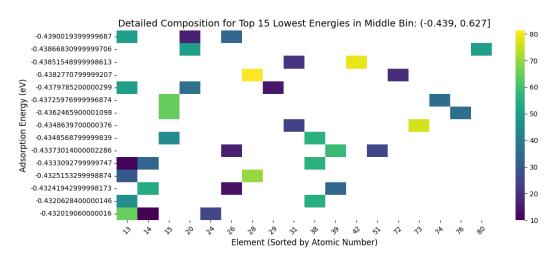


Figure 6: Composition heatmap for specific range of adsorption energies on shell 2.

390 F Diffusion model details

391 Forward Noising Process

Let z be the initial input x_0 for diffusion, we define a fixed forward (noising) kernel that does not depend on y:

$$q(x_{1:T} \mid \mathbf{z}) = \prod_{t=1}^{T} q(x_t \mid x_{t-1}), \quad q(x_t \mid x_{t-1}) = \mathcal{N}\left(x_t; \sqrt{1 - \beta_t} x_{t-1}, \beta_t \mathbf{I}\right)$$

$$x_T \sim \mathcal{N}(0, \mathbf{I}).$$

Here, x_T is the noisy input at timestep T, drawn from a multivariate normal distribution with zero mean and identity covariance matrix (I). This x_T will undergo a series of denoising steps to gradually transform into an idea latent embedding on the given information using the formula given above. Marginally,

$$q(x_t \mid \mathbf{z}) = \mathcal{N}(x_t; \sqrt{\bar{\alpha}_t} \mathbf{z}, (1 - \bar{\alpha}_t) \mathbf{I}), \quad \bar{\alpha}_t = \prod_{s=1}^t (1 - \beta_s)$$

Learned Reverse (Denoising) Process To reverse the forward diffusion process, we train a neural network ϵ_{θ} to predict the noise added to a noised sample x_t at time step t. This prediction is conditioned not only on x_t and t, but also on additional conditioning variables specific to the task:

$$\epsilon_{\theta} = \epsilon_{\theta}(x_t, t, \mathbf{s}, \mathbf{e}, \mathbf{l})$$

- where s, e and l denote external conditioning inputs, composed by three counterparts as language description for elements and properties, elements in categories and properties in numerical value.
- Based on this predicted noise, we define the reverse transition distribution as a Gaussian:

$$p_{\theta}(x_{t-1} \mid x_t, \mathbf{s}, \mathbf{e}, \mathbf{l}) = \mathcal{N}(x_{t-1}; \mu_{\theta}(x_t, t, \mathbf{s}, \mathbf{e}, \mathbf{l}), \Sigma_t)$$

Here, the mean μ_{θ} and the variance Σ_t are typically parameterized as follows:

$$\Sigma_t = \beta_t \mathbf{I}, \quad \mu_{\theta}(x_t, t, \mathbf{s}, \mathbf{e}, \mathbf{l}) = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \cdot \epsilon_{\theta}(x_t, t, \mathbf{s}, \mathbf{e}, \mathbf{l}) \right)$$

- This formulation follows the DDPM framework, where ϵ_{θ} learns to approximate the true noise, and the predicted mean guides the reverse sampling trajectory toward denoised data **z**.
- 407 Training Objective
- We still minimize the simple noise prediction loss, but conditioned on (s, e, l):

$$\mathcal{L}(\theta) = \mathbb{E}_{\mathbf{z}, \mathbf{s}, \mathbf{e}, \mathbf{l}, \epsilon \sim \mathcal{N}(0, I), t} \left[\left\| \epsilon - \epsilon_{\theta} \left(\sqrt{\bar{\alpha}} \mathbf{z} + \sqrt{1 - \bar{\alpha}_{t}} \epsilon, t, \mathbf{s}, \mathbf{e}, \mathbf{l} \right) \right\|^{2} \right].$$

- Here, \mathbf{s} , \mathbf{e} , \mathbf{l} are drawn from their respective empirical distributions, ϵ is standard Gaussian noise, and t is uniform on $\{1, \dots, T\}$.
- Inference (Sampling) At inference time, we begin by drawing a pure noise sample x_T from the standard normal distribution $\mathcal{N}(0,\mathbf{I})$. Then, for each timestep t counting down from T to 1, we denoise by sampling

$$x_{t-1} \sim \mathcal{N}\left(\mu_{\theta}\left(x_{t}, t, \mathbf{s}, \mathbf{e}, \mathbf{l}\right), \beta_{t} \mathbf{I}\right)$$

where μ_{θ} is the predicted mean of our network conditioned on the current noisy state x_t , the timestep t, and the three conditioning signals s (adsorbate), e (adsorption energy) and ℓ (language description) which will be transferred as vectors \mathbf{s} , \mathbf{e} and \mathbf{l} correspondingly. As we step backward through time, the noise magnitude β_t shrinks according to our predefined schedule, gradually transforming the initial noise into our target latent sample. After completing the final step at t=1, the resulting \mathbf{z} is returned as a sample approximately drawn from the desired conditional distribution p_{data} (\mathbf{z} | \mathbf{s} , \mathbf{e} , \mathbf{l}).

Table 3: CIAE pretraining reconstruction performance. The evaluation includes test sets from both the pretraining domain dataset (OC20) and the out-of-domain dataset (HEA). Metrics reported are **Mean Square Error (MSE)** and **Cosine Similarity (Sim)**.

Dataset	Avg. MSE \downarrow	Avg. Sim (shell 1) \uparrow	Avg. Sim (shell 2) ↑	Avg. Sim (all shells) ↑
OC20 test set	3.4974	0.8767	0.9935	0.9351
HEA test set	5.1630	0.9785	0.9471	0.9628

420 G More details about experiments

421 G.1 Baseline.

We construct a chemically agnostic baseline that bypasses the CIAE and feeds an unstructured, image-based representatio directly to the DiT model. Starting from a composition vector $\mathbf{c} \in \mathbb{R}^{118}$, we zero-pad it to an augmented vector $\mathbf{c}' \in \mathbb{R}^{120}$ for reshaping compatibility, where $\mathbf{c}' = \mathrm{pad}(\mathbf{c}; (0, 2))$; \mathbf{c}' is then reshaped into a matrix $\mathbf{M} = reshape(\mathbf{c}', (15, 8)) \in \mathbb{R}^{15 \times 8}$, mapping the elemental features into a rudimentary spatial grid devoid of any embedded chemical priors. Finally, to create a standardized input, this matrix \mathbf{M} is padded into a larger zero-tensor $\mathbf{X} \in \mathbb{R}^{48 \times 48}$, which is defined as $X_{i,j} = M_{i,j}$ if $1 \le i \le 15$, $1 \le j \le 8$ and 0 otherwise. This final tensor \mathbf{X} serves as the direct input to the diffusion model.

430 G.2 Metrics

431

432

433

434

435

437

438

439

440

441

442

443

444

445

• Fréchet Distance (FD). Motivated by image generation evaluation (Heusel et al., 2017), we adopt the Fréchet distance for comparing the generated distribution and input distribution of surface compositions. We use CIAE encoder to project both real (X_r) and generated (X_g) compositions into the same chemical-aware latent space. The FD is then calculated based on the means (μ_r, μ_g) and covariances (Σ_r, Σ_g) of the latent embeddings:

$$\mathrm{FD} = ||\boldsymbol{\mu}_r - \boldsymbol{\mu}_g||_2^2 + \mathrm{Tr} \left(\boldsymbol{\Sigma}_r + \boldsymbol{\Sigma}_g - 2(\boldsymbol{\Sigma}_r \boldsymbol{\Sigma}_g)^\frac{1}{2}\right)$$

Lower FD scores indicate a closer match between the real and generated distributions.

• **Leakage.** For a ground truth composition c, define its support (allowed set of elements) as $S = \{i \mid c_i > 0\}$. Given a generated composition \hat{c} , leakage quantifies the sum of all proportions fall outside the support in the disallowed set S^c . For a batch of N samples, we define:

$$\text{Leakage} = \frac{1}{N} \sum_{j=1}^{N} \sum_{i \in S_j^c} \hat{c}_i^{(j)}$$

where S_j^c is the disallowed set for the j-th ground truth sample $\mathbf{c}^{(j)}$. Lower values indicate better consistency of generated compositions.

• Cosine Similarity (Sim). It provides a direct point-wise measure of generation quality by evaluating the alignment between a pair of ground truth vector c and generated vector ĉ. Sim computes the average cosine similarity over N samples:

$$Sim = \frac{1}{N} \sum_{j=1}^{N} \frac{\mathbf{c}^{(j)} \cdot \hat{\mathbf{c}}^{(j)}}{||\mathbf{c}^{(j)}||_{2} ||\hat{\mathbf{c}}^{(j)}||_{2}}$$

where the \cdot sign denotes the dot product. Higher values indicate higher similarity.

G.3 CIAE pretraining results

These results in Table 3 show that CIAE can efficiently map the sparse inputs to the chemically consistent latent space and achieve low reconstruction error and high alignment with target elemental compositions across both shells, with reasonable transferability to unseen HEA dataset.

H Limitations and Future Work

Looking ahead, we plan to redesign the current PTR-initialized VGG encoder with a purpose-built 451 CIAE tailored especially for catalytic surface compositions, exploring novel architectures (e.g., axial attention) and conducting broader ablation studies. For stage I pretraining, we will move 453 beyond generic classification pretraining task toward objectives that better align the latent space 454 with representation of surface compositions. We will also enrich the multi-physical conditions with 455 obtainable surface information related with ensemble and ligand effects, such as the active site motifs, 456 facet labels, and simple geometric details such as distances from the adsorbate center to first shell 457 atoms, in order to develop more expressive embeddings for these inputs. Together, these upgrades aim 458 to yield more uniquely determined generations and a stronger, end-to-end path from target properties 459 to deployable catalyst candidates.