Supplementary Information

Dismai-Bench: Benchmarking and designing generative models using disordered materials and interfaces

Adrian Xiao Bin Yong[*1, 2], Tianyu Su[1, 2], and Elif Ertekin[*2, 3]

[1]Department of Materials Science and Engineering, University of Illinois Urbana-Champaign, 1304 West Green Street, Urbana 61801 IL, USA

[2]Materials Research Laboratory, University of Illinois Urbana-Champaign, 104 South Goodwin Avenue, Urbana 61801 IL, USA

[3]Department of Mechanical Science and Engineering, University of Illinois Urbana-Champaign, 1206 West Green Street, Urbana 61801 IL, USA

[*]Corresponding authors; Emails: axyong2@illinois.edu, ertekin@illinois.edu

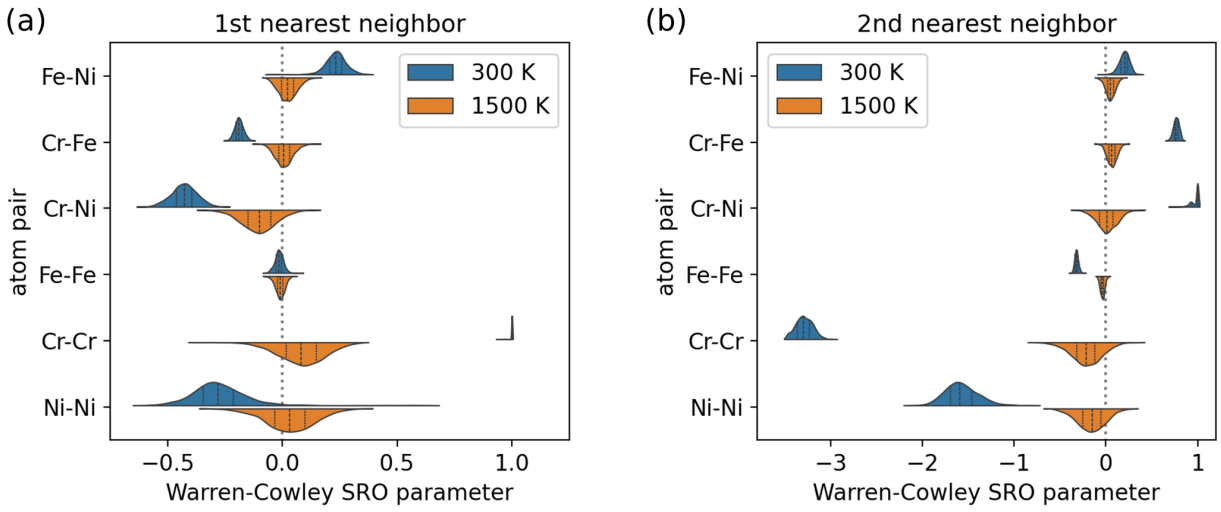

Figure S1: Warren-Cowley short-range order (SRO) parameter distributions of 300 K and 1500 K alloy training datasets (no SRO filtering). The SRO distributions are shown for the (a) 1st and (b) 2nd nearest neighbor interactions.

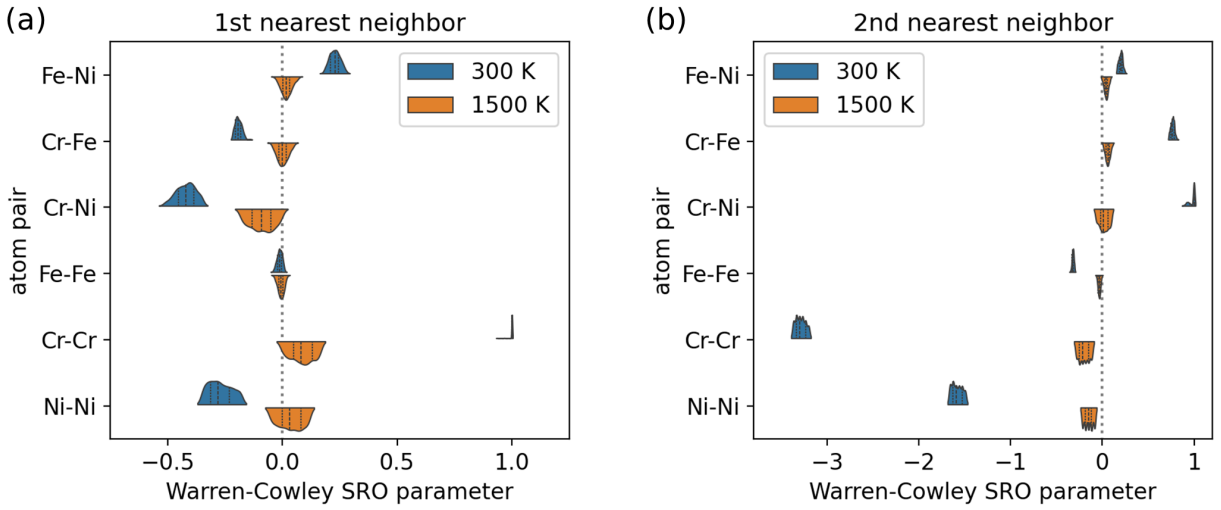

Figure S2: Warren-Cowley SRO parameter distributions of 300 K and 1500 K alloy training datasets (with SRO filtering). The SRO distributions are shown for the (a) 1st and (b) 2nd nearest neighbor interactions.

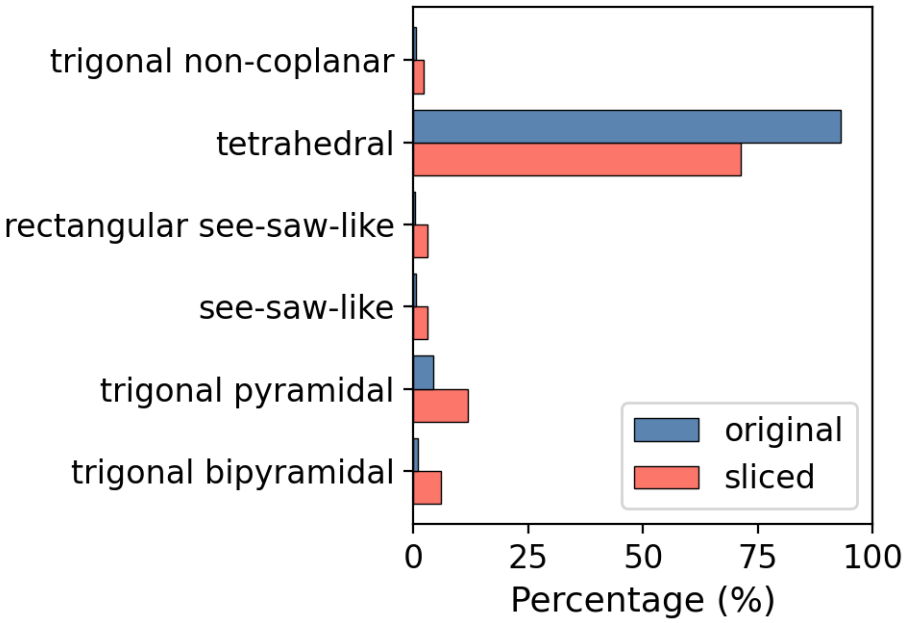

Figure S3: Coordination motif distribution of amorphous silicon, showing the comparison between the original 100,000-atom structure and the sliced 256-atom structures.

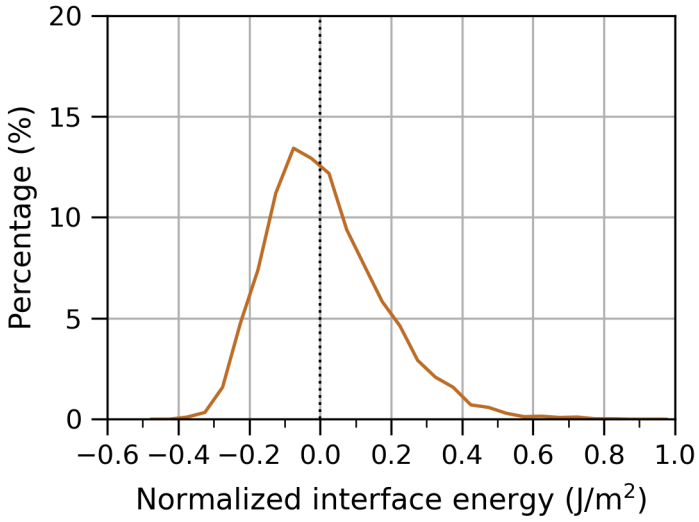

Figure S4: Normalized interface energy distribution of randomly generated interface structures. All structures were relaxed using M3GNet.

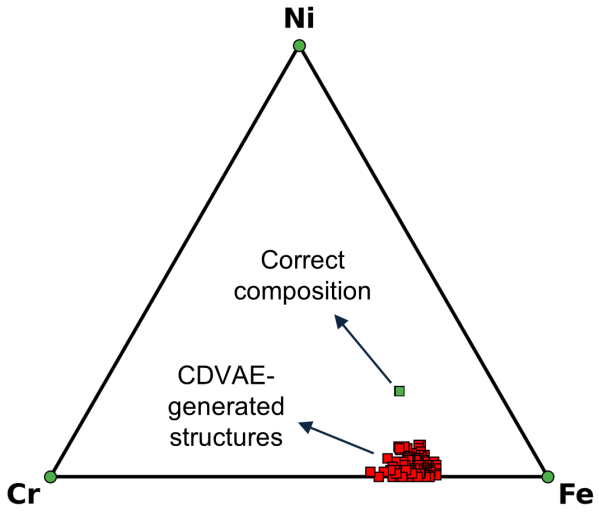

Figure S5: Phase diagram of the alloy structures generated by CDVAE (with atomic species denoising). The correct composition, $Fe_{60}Ni_{20}Cr_{20}$, is indicated by the green square. The composition of 100 CDVAE-generated structures are indicated by the red squares.

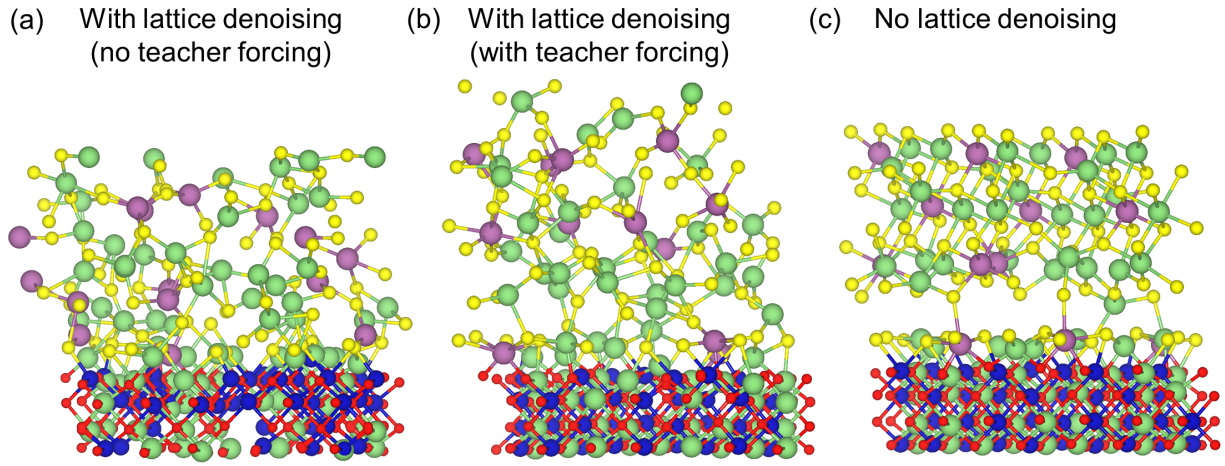

Figure S6: Example DiffCSP-generated interface structures (a) with lattice denoising but no teacher forcing, (b) with lattice denoising and teacher forcing, and (c) without lattice denoising.

We trained DiffCSP models with lattice denoising (and teacher forcing) to compare with DiffCSP models without lattice denoising. For the disordered interface and alloy datasets, teacher forcing was applied for the first 150 epochs. For the amorphous Si dataset, teacher forcing was applied for the first 500 epochs. The weight of the lattice cost was set to 1 for all models. All other hyperparameters were set to be the same as the DIffCSP models without lattice denoising.

Table S1: Dismai-Bench metrics for DiffCSP with lattice denoising and teacher forcing. Each metric is represented by the average value over 3 separately trained models. The minimum and maximum values are shown in round brackets. The difference in the metric value between DiffCSP with lattice denoising and without lattice denoising is shown in square brackets. DiffCSP performs worse with lattice denoising than without lattice denoising.

Disordered interface

d_{Li} (min, max)	d_{Co} (min, max)	d_{Sc} (min, max)	d_{all} (min, max)	% struc failed (min, max)
0.215 [**+0.171**] (0.0853, 0.297)	0.273 [**+0.245**] (0.0905, 0.420)	0.317 [**+0.222**] (0.240, 0.407)	0.192 [**+0.155**] (0.0671, 0.278)	43.8 [**+36.7**] (11.9, 79.3)

Amorphous Si

d_{motif} (min, max)	d_{rdf} (min, max)	d_{angle} (min, max)	% struc failed (min, max)
0.863 [**+0.798**] (0.405, 1.11)	9.70 [**+8.31**] (8.70, 10.6)	0.0492 [**+0.0390**] (0.0260, 0.0621)	8.7 [**+8.7**] (6.0, 12.6)

Alloy (300 K, narrow SRO)

d_{cluster} (min, max)	% struc w/ vac (min, max)	% struc failed (min, max)
0.0824 [**+0.0179**] (0.0621, 0.102)	99.5 [**+5.13**] (98.4, 100)	0.00 [**+0.00**] (0.00, 0.00)

Table S2: M3GNet relaxation results of disordered interfaces generated by the generative models. n_{steps} is the number of M3GNet relaxation steps required to relax the generated structures. E_{initial} and E_{final} are the M3GNet-calculated energies of the unrelaxed and relaxed structures respectively. The mean and standard deviation for each quantity are listed. For each architecture, 3 separate models were trained and the results were averaged.

Model	n_{steps}		E_{initial} - E_{final} (meV/atom)	
	mean	std	mean	std
CDVAE	10.3	8.71	3.30	1.46
DiffCSP	68.2	53.9	12.3	9.68
CrysTens	171	74.3	640	137
UniMat	257	116	989	70.0
CryinGAN	147	72.4	234	67.3

Table S3: SOAP-GAP relaxation results of amorphous Si generated by CDVAE and DiffCSP. n_{steps} is the number of SOAP-GAP relaxation steps required to relax the generated structures. E_{initial} and E_{final} are the SOAP-GAP-calculated energies of the unrelaxed and relaxed structures respectively. The mean and standard deviation for each quantity are listed. For each architecture, 3 separate models were trained and the results were averaged.

Model	n_{steps}		E_{initial} - E_{final} (meV/atom)	
	mean	std	mean	std
CDVAE	22.5	16.3	8.64	45.3
DiffCSP	140	39.0	275	21.6

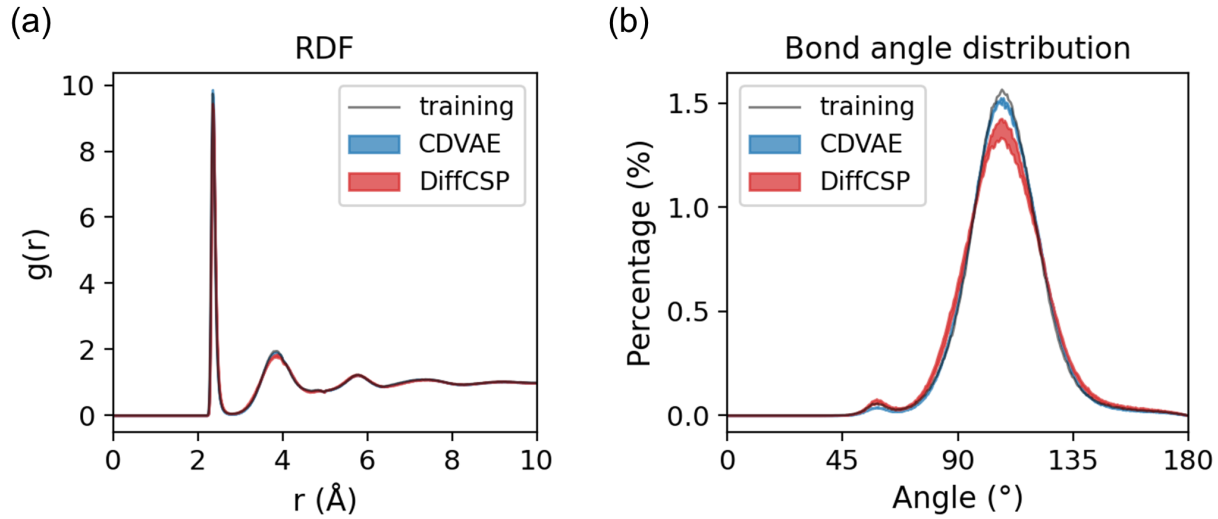

Figure S7: (a) Radial distribution functions (RDFs) and (b) bond angle distributions of amorphous Si structures generated by CDVAE and DiffCSP. The distributions of the training dataset are also shown for reference.

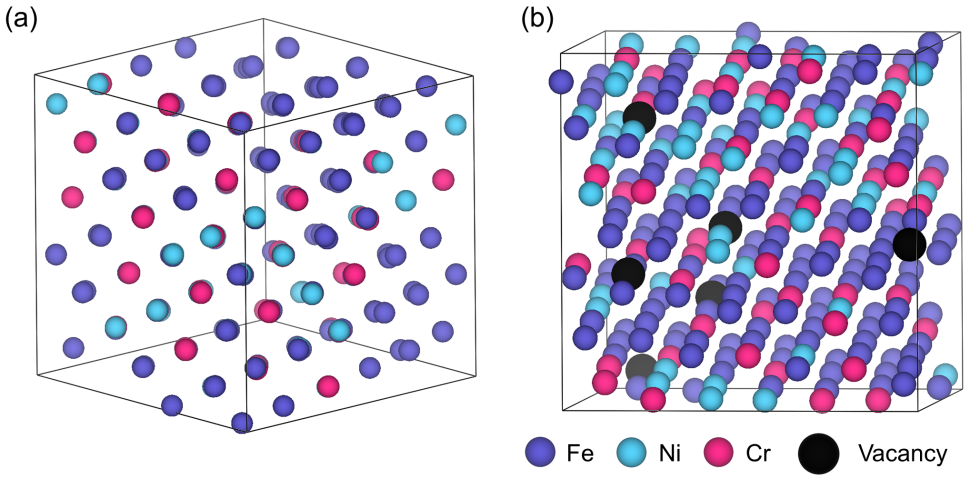

Figure S8: Examples of CDVAE-generated alloy structures with (a) noisy lattice and (b) vacancies.

CDVAE

(a) 300 K, wide SRO:

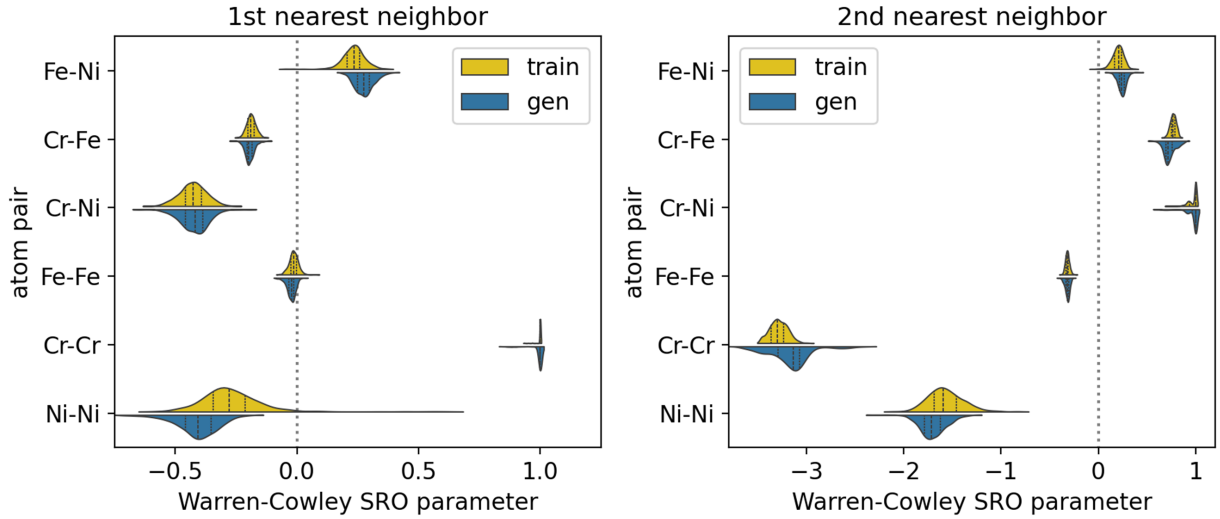

(b) 1500 K, wide SRO:

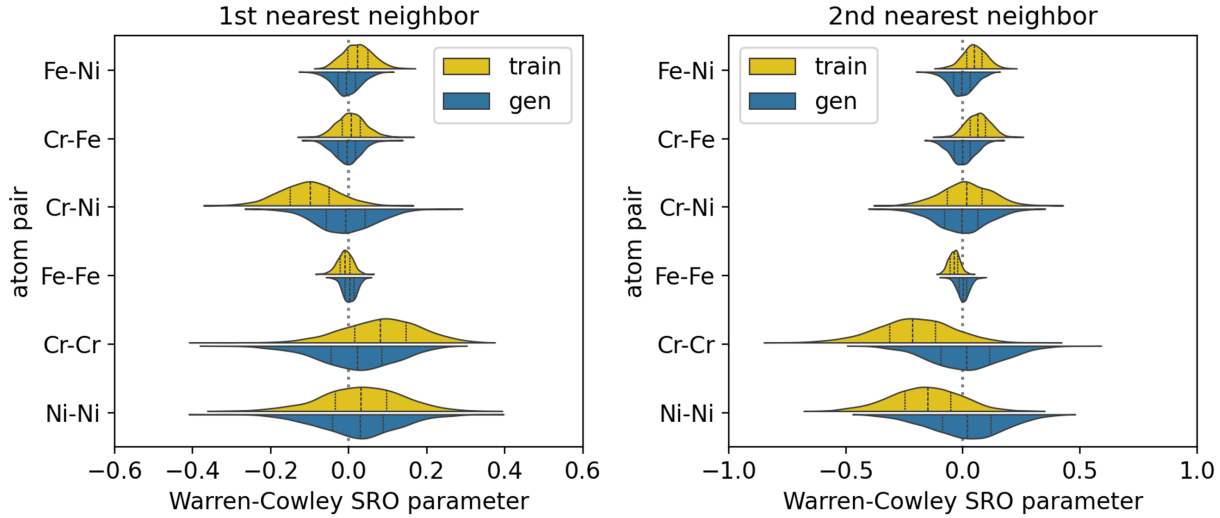

Figure S9: Warren-Cowley SRO parameter distributions of (a) 300 K and (b) 1500 K alloy structures (wide SRO) generated by CDVAE. The SRO distributions are shown for the 1st and 2nd nearest neighbor interactions.

DiffCSP

(a) 300 K, wide SRO:

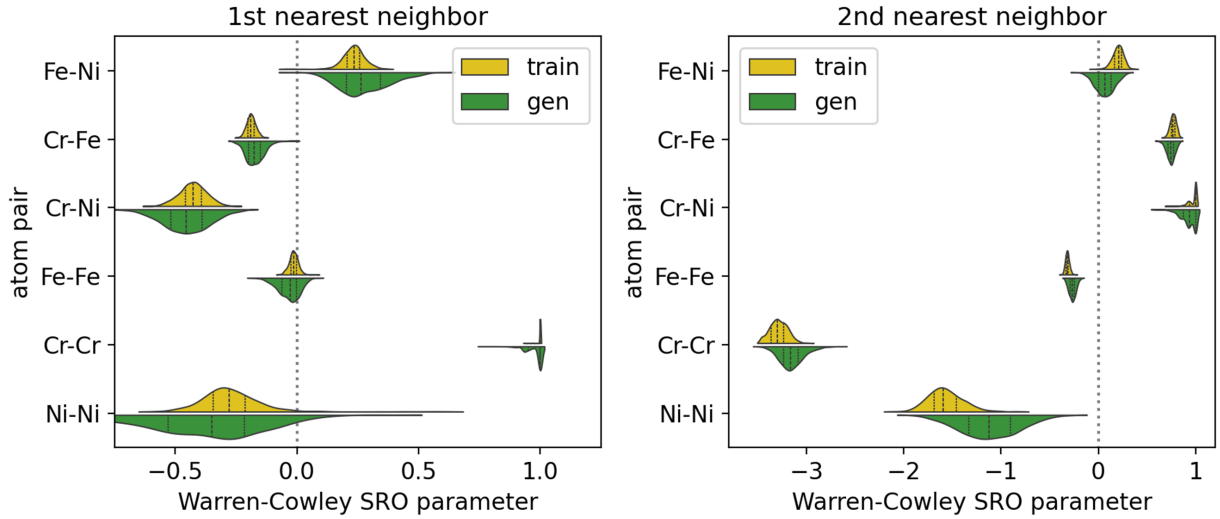

(b) 1500 K, wide SRO:

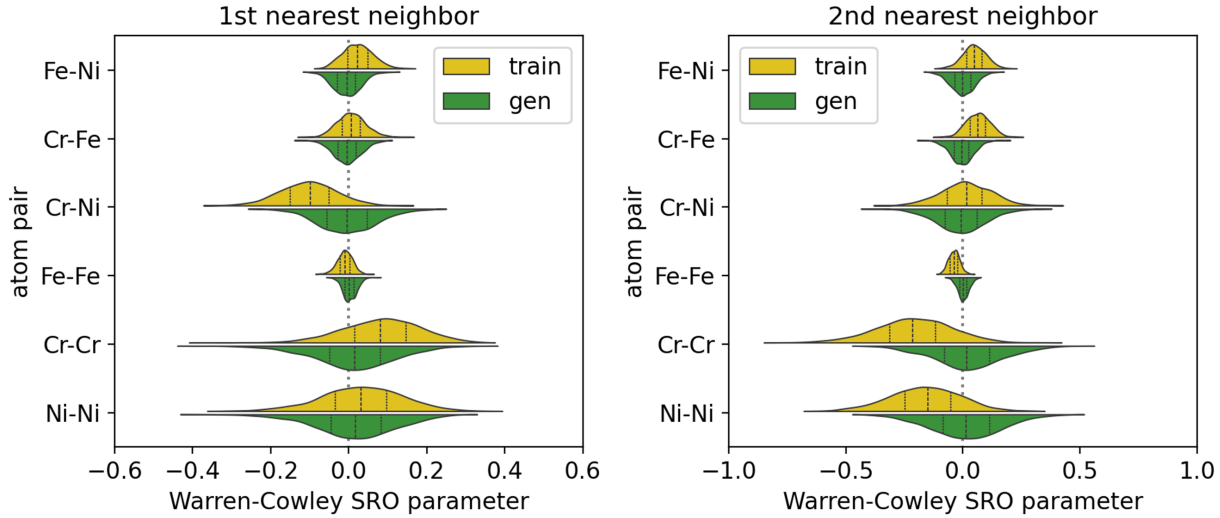

Figure S10: Warren-Cowley SRO parameter distributions of (a) 300 K and (b) 1500 K alloy structures (wide SRO) generated by DiffCSP. The SRO distributions are shown for the 1st and 2nd nearest neighbor interactions.

CryinGAN

(a) 300 K, wide SRO:

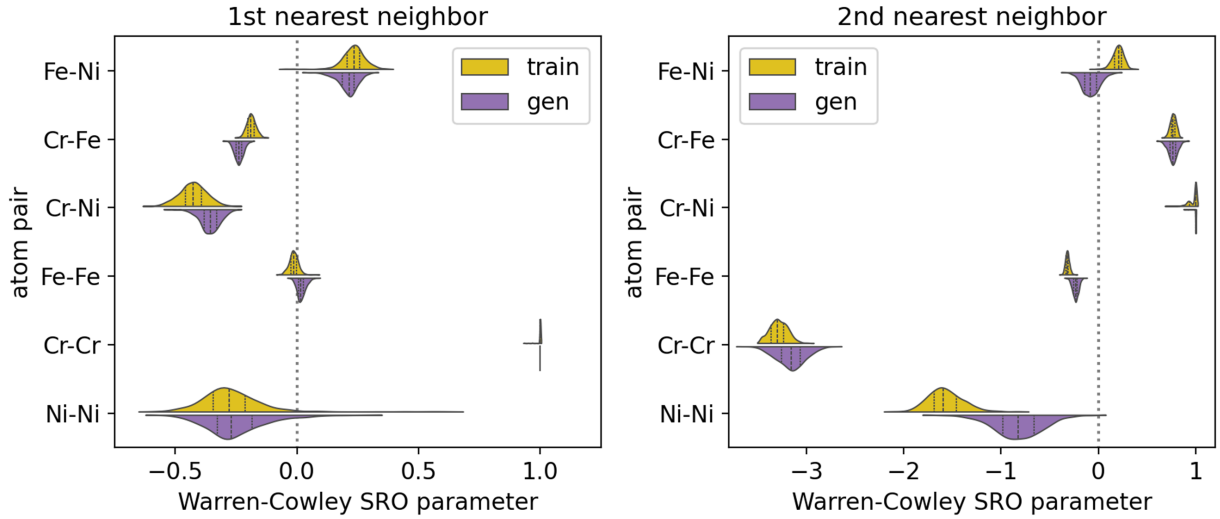

(b) 1500 K, wide SRO:

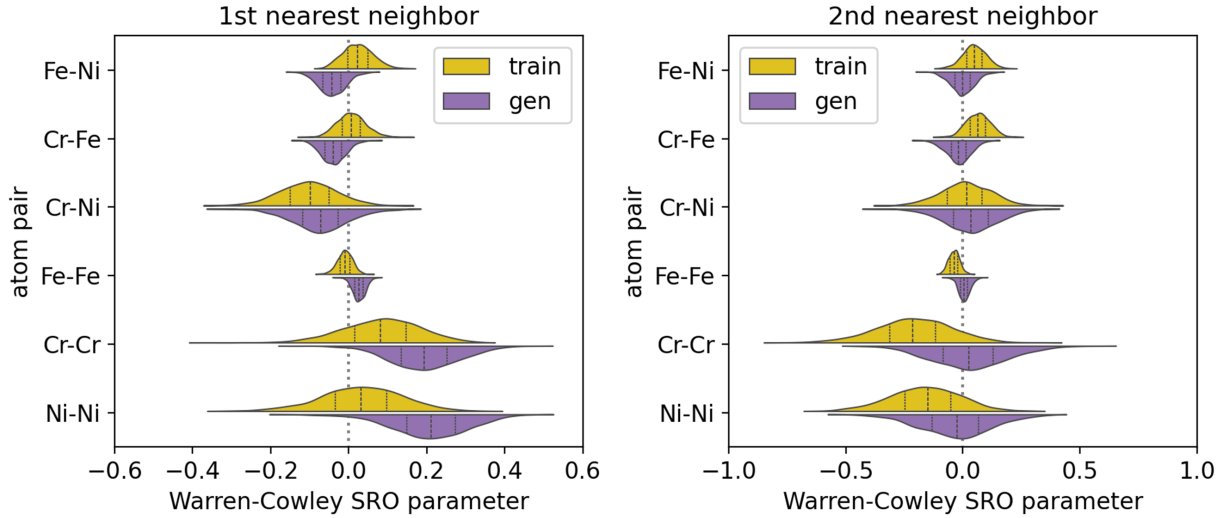

Figure S11: Warren-Cowley SRO parameter distributions of (a) 300 K and (b) 1500 K alloy structures (wide SRO) generated by CryinGAN. The SRO distributions are shown for the 1st and 2nd nearest neighbor interactions.

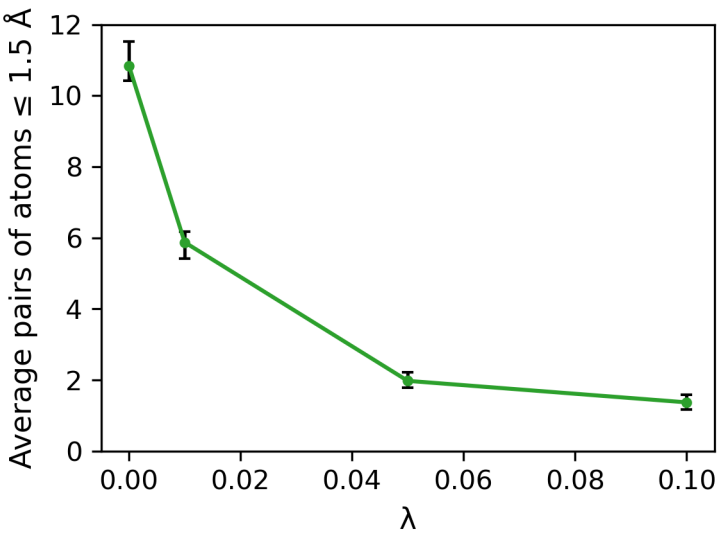

Figure S12: Average number of unique pairs of atoms with bond distance $\leq$ 1.5 Å for interface structures generated using CryinGAN trained with different λ values. As λ increases, fewer atoms are generated too close to each other. For each λ value, 3 separate models were trained, 1000 structures were generated using each model, and the counts were averaged across the models.

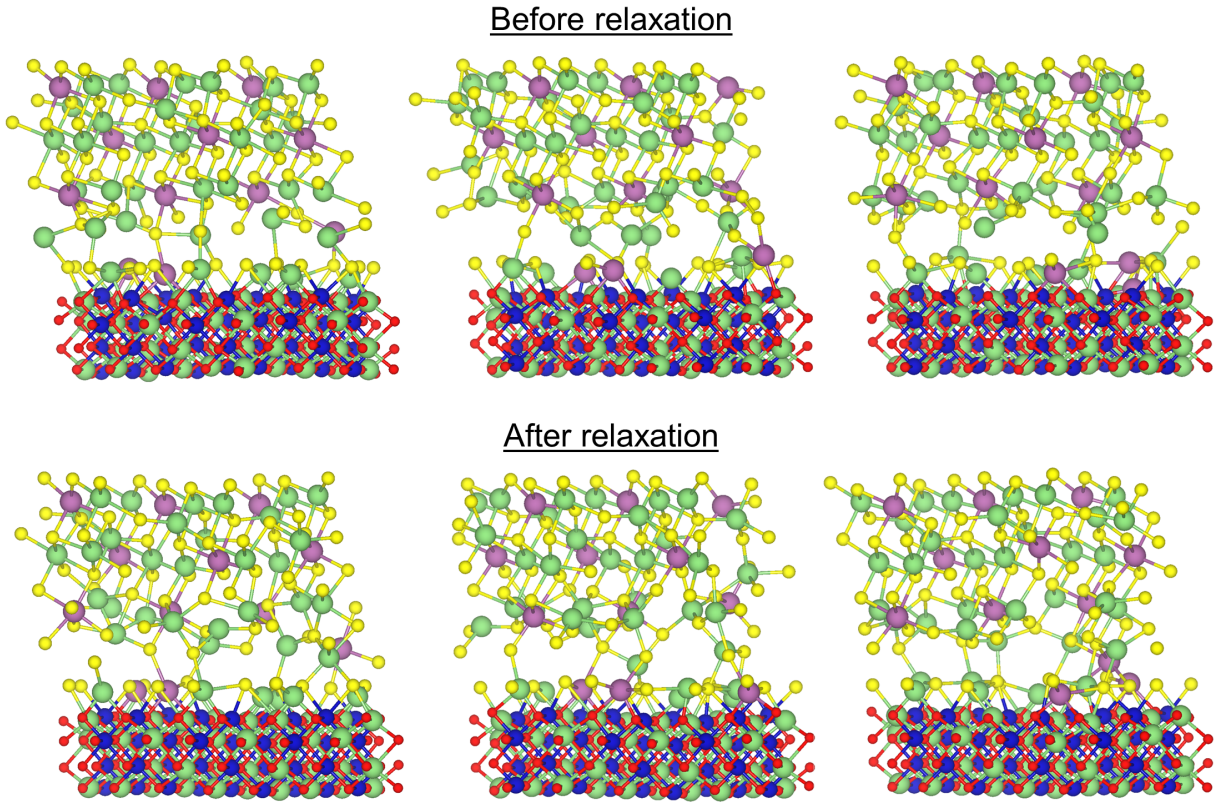

Figure S13: Examples of CryinGAN-generated structures before relaxation (as-generated) and after M3GNet relaxation.

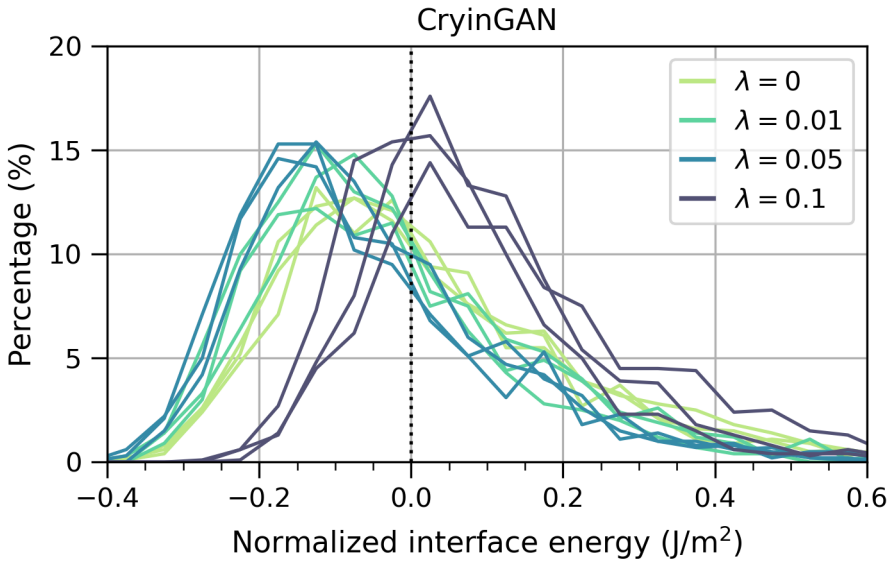

Figure S14: Normalized interface energy distributions of structures generated using CryinGAN trained with different λ values (0, 0.01, 0.05, and 0.1). For each λ value, 3 separate models were trained. The generated structures were relaxed using M3GNet, and the interface energies shown are based on M3GNet-calculated energies.

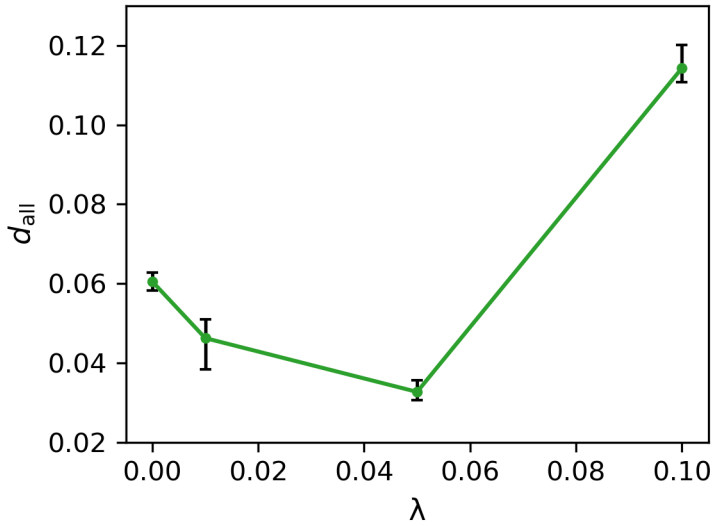

Figure S15: Euclidean distance of the coordination motif fingerprint of all cations for interface structures generated using CryinGAN trained with different λ values.

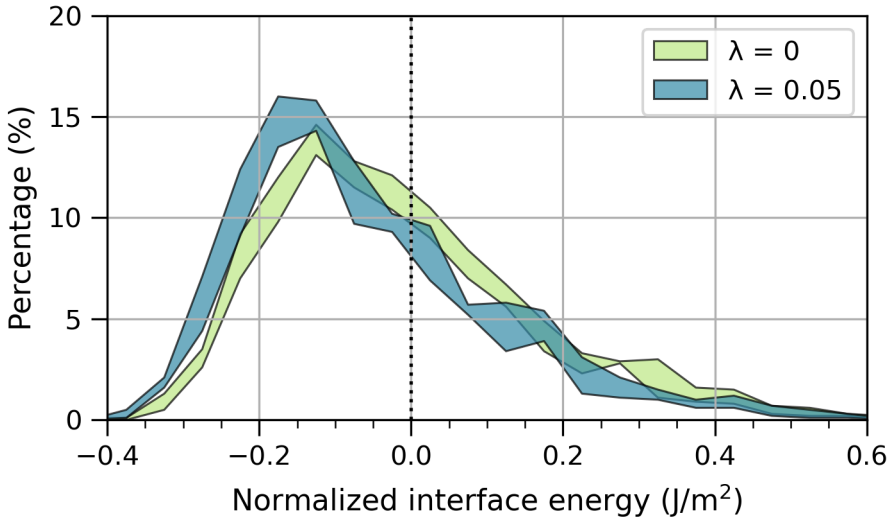

Figure S16: Normalized interface energy distributions of structures generated using CryinGAN trained with $\lambda = 0$ (200,000 epochs) and $\lambda = 0.05$ (100,000 epochs). The generated structures were relaxed using M3GNet, and the interface energies shown are based on M3GNet-calculated energies. CryinGAN runs approximately two times faster when the bond distance discriminator is not used ($\lambda = 0$), but the use of the bond distance discriminator still yields a lower interface energy distribution for the same amount of training time.

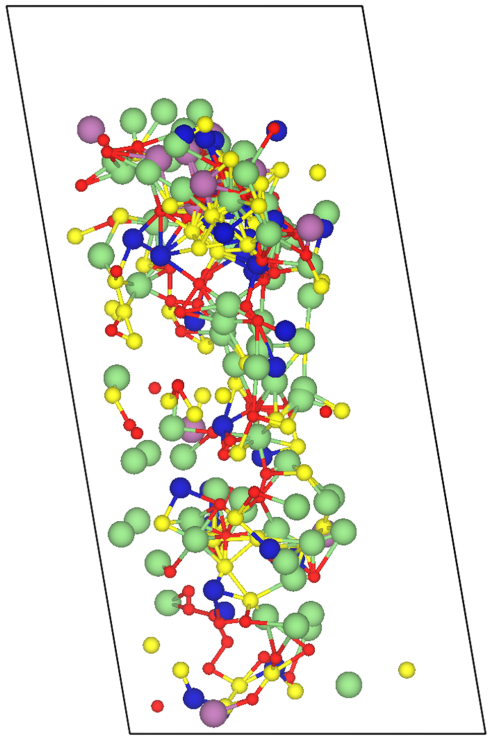

Figure S17: Example generated structure of a GAN model using Crystal Graph Convolutional Neural Networks (CGCNN) as the discriminator. The GAN was unable to learn to generate meaningful interface structures.

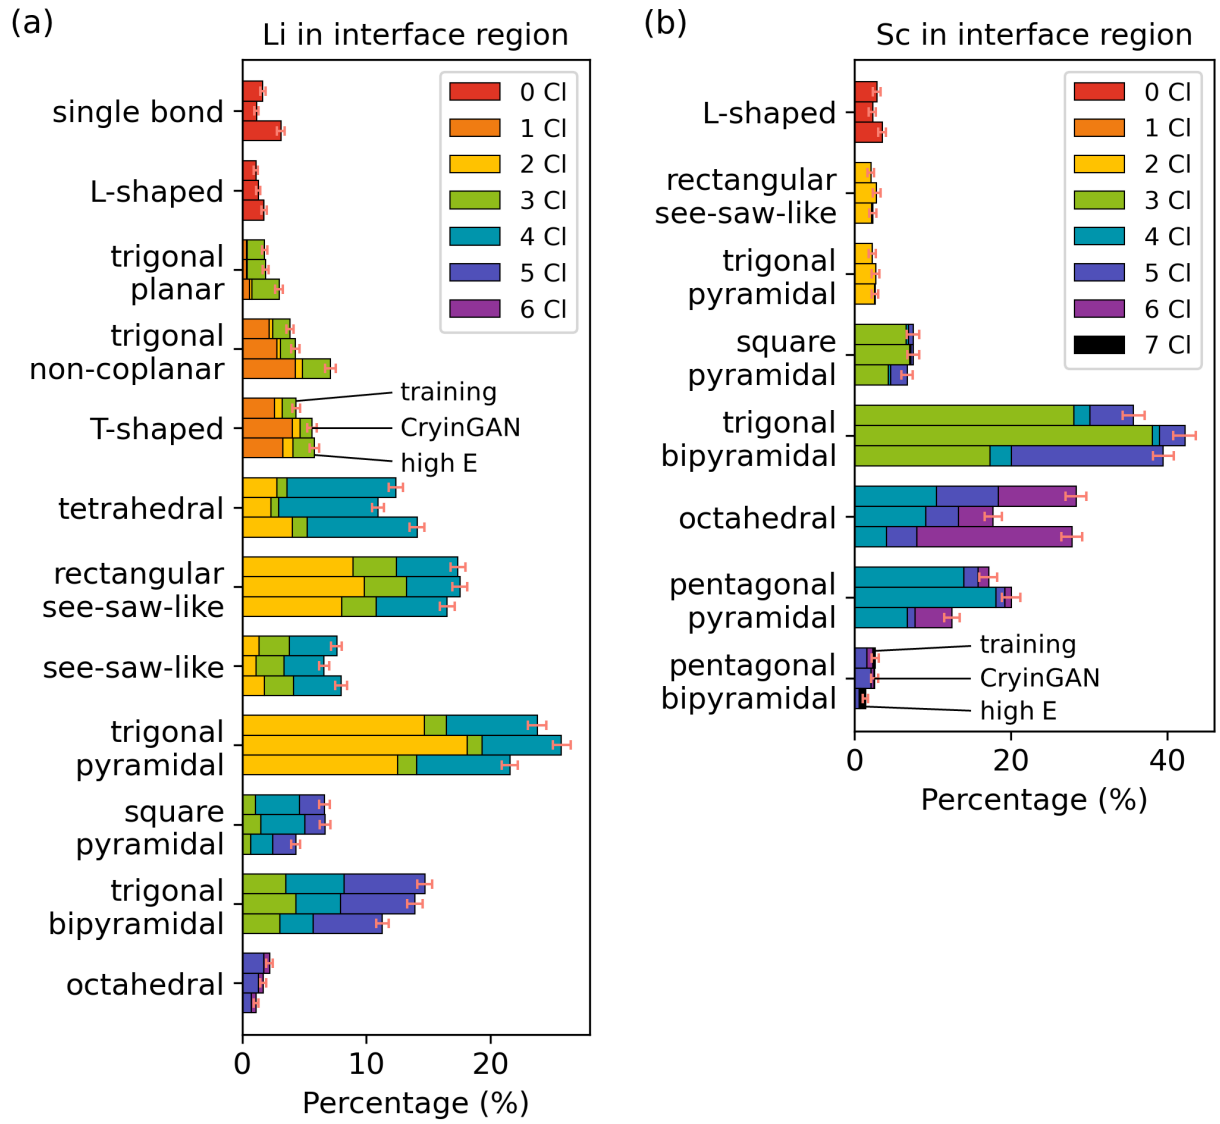

Figure S18: Coordination motif distributions of (a) Li and (b) Sc in the interface region, where each coordination motif is further subdivided based on the number of Cl bonds. The distributions of three datasets are shown: (1) training structures with low interface energy, (2) CryinGAN-generated structures, and (3) structures with high interface energy. All structures were relaxed using M3GNet followed by DFT calculations. Error bars represent 95 % bootstrap confidence intervals.

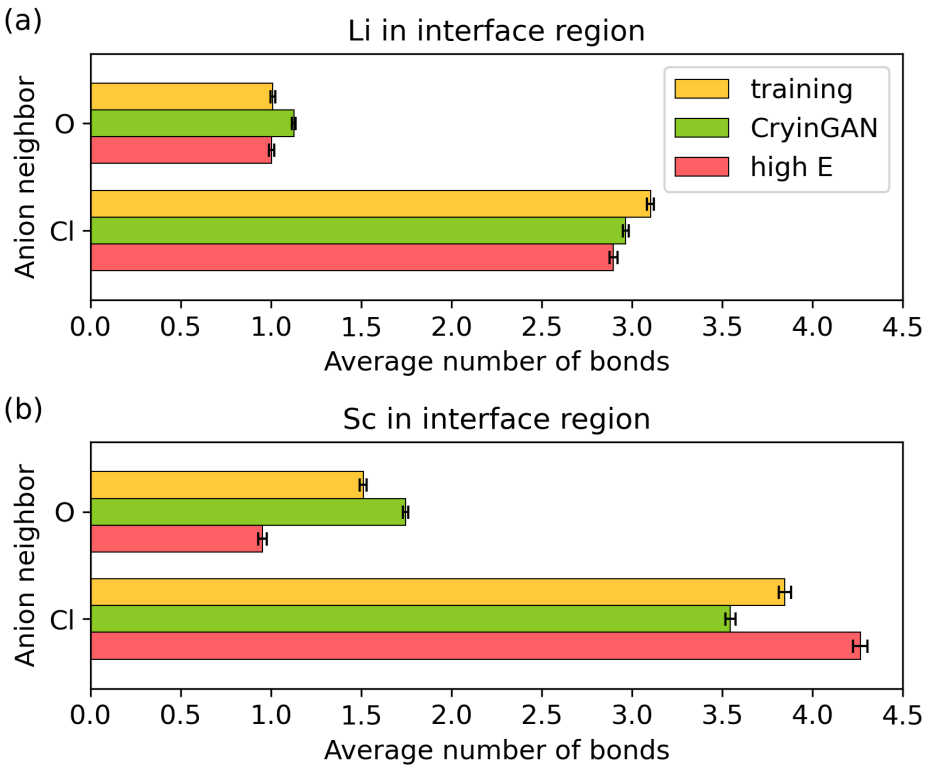

Figure S19: Histograms of the average number of O and Cl bonds for (a) Li and (b) Sc in the interface region. Error bars represent 95 % bootstrap confidence intervals.

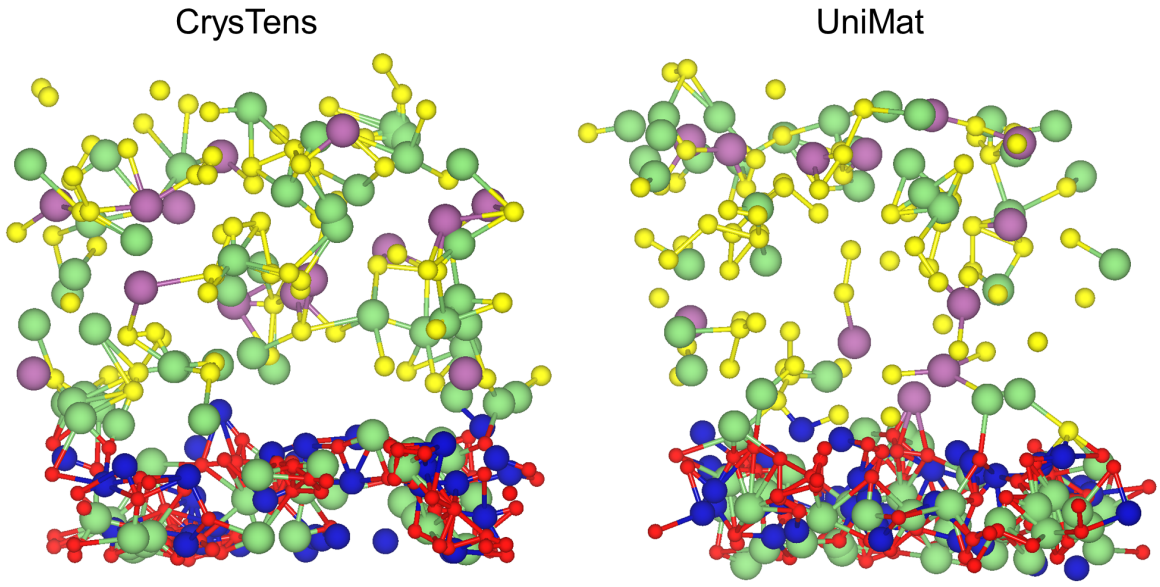

Figure S20: Example interface structures generated by CrysTens and UniMat when trained on structures with randomized atom orderings.

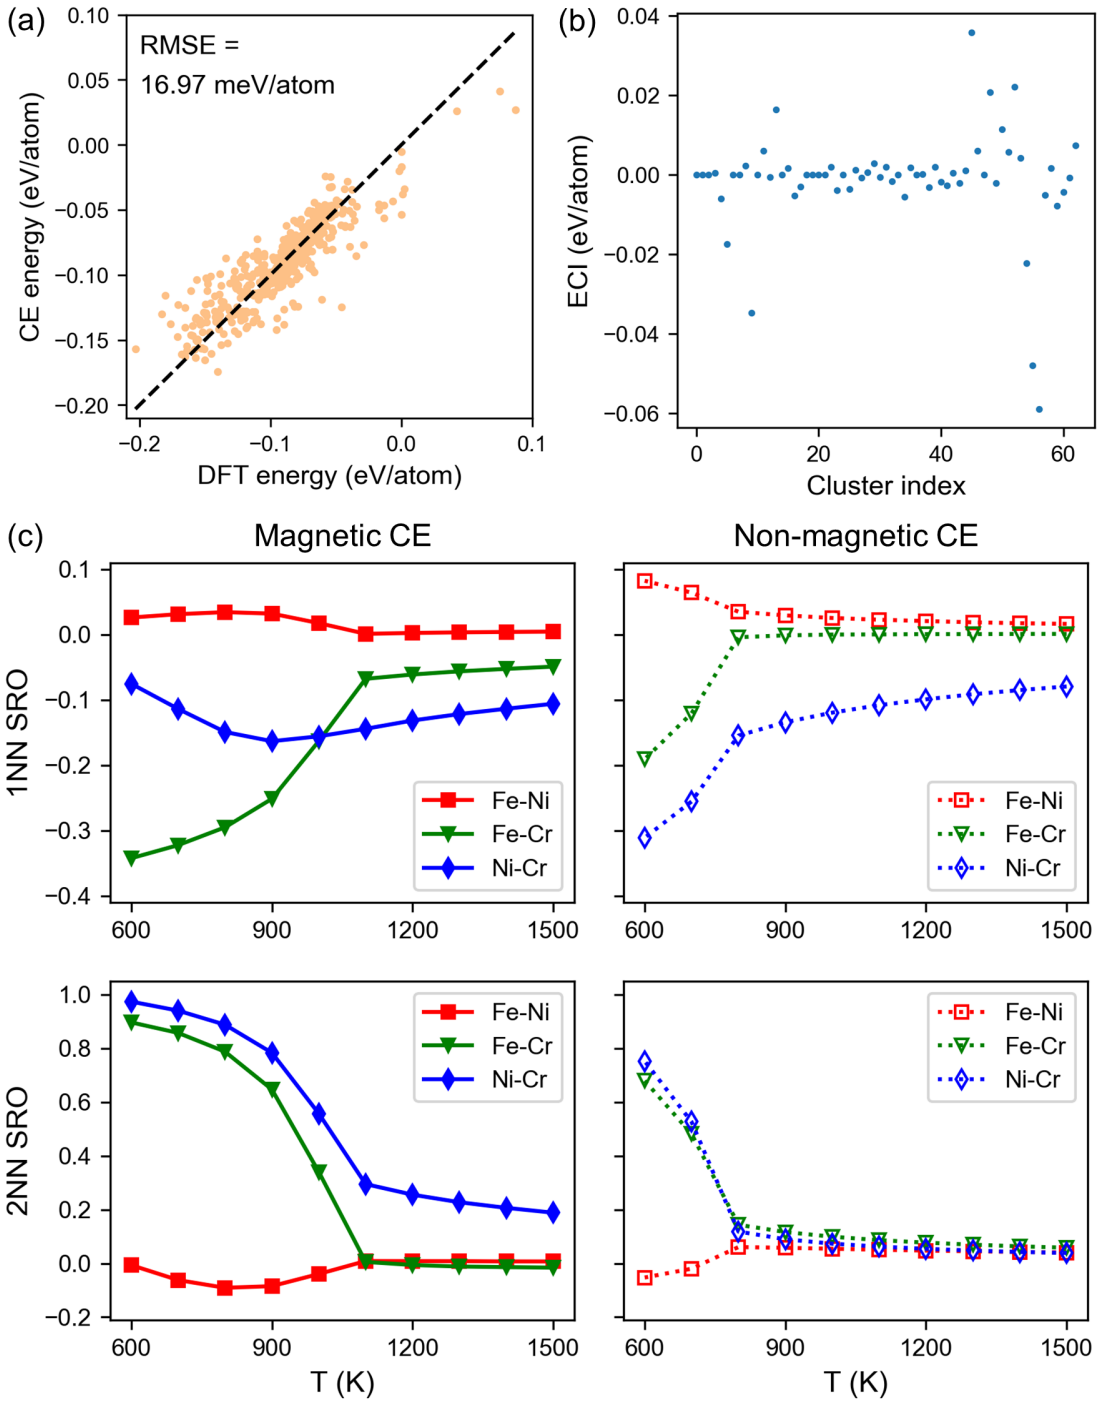

Figure S21: (a) Plot of cluster expansion (CE) energies against DFT energies. The root mean square error (RMSE) of the CE fit is 16.97 meV/atom. (b) All effective cluster interactions (ECI) values of the CE fit. (c) Plots of the Warren-Cowley SRO parameter against temperature. Plots on the left are adapted from ref. 1 and correspond to an alloy with composition $Fe_{56}Cr_{21}Ni_{23}$ and 4000 atoms. Plots on the right correspond to the alloy of this work, with composition $Fe_{60}Ni_{20}Cr_{20}$ and 256 atoms. Plots on the left use the full CE model with magnetic terms as described in ref. 1, whereas plots on the right use only terms corresponding to the seven chemical dimers (non-magnetic). The top and bottom plots correspond to the 1st and 2nd nearest neighbor interactions respectively. The Dismai-Bench alloy shows qualitatively similar SRO trends to the larger magnetic alloy.

Table S4: Configurations of the interfaces calculated. The slab orientation, number of layers, and average lattice mismatch between any given two slabs are listed. M represents the transformation matrix used to transform the lattice vectors of a given slab surface $(\vec{u}, \vec{v})$ into the superlattice vectors of the interface $(\vec{u}_s, \vec{v}_s)$, according to the relation $(\vec{u}_s, \vec{v}_s) = M \cdot (\vec{u}, \vec{v})$. Note that the number of layers of the $Li_3ScCl_6(100)$ slab is indicated for the randomly generated $LiCoO_2(110)$-$Li_3ScCl_6(100)$ structures, and does not include the atoms randomly generated in the interface region.

Slab 1	number of layers	M	Slab 2	number of layers	M	average mismatch (%)
$LiCoO_2(110)$	4	$\begin{pmatrix} 1 & 2 \\ -3 & 3 \end{pmatrix}$	$Li_3ScCl_6(100)$	9	$\begin{pmatrix} 1 & 1 \\ -2 & 1 \end{pmatrix}$	2.17
$Li_2O(100)$	9	$\begin{pmatrix} 1 & 2 \\ -4 & 2 \end{pmatrix}$	$LiCl(100)$	4	$\begin{pmatrix} 2 & 0 \\ 0 & 4 \end{pmatrix}$	0.545
$Li_2O(110)$	4	$\begin{pmatrix} 1 & 0 \\ 0 & 3 \end{pmatrix}$	$Li(100)$	6	$\begin{pmatrix} 1 & 0 \\ 0 & 3 \end{pmatrix}$	3.00
$Li_2O(110)$	4	$\begin{pmatrix} -2 & 1 \\ 3 & 2 \end{pmatrix}$	$LiCl(100)$	4	$\begin{pmatrix} 2 & 1 \\ -2 & 3 \end{pmatrix}$	2.13
$Li_2O(110)$	4	$\begin{pmatrix} -2 & 1 \\ 3 & 2 \end{pmatrix}$	$MgS(100)$	4	$\begin{pmatrix} 2 & 1 \\ -2 & 3 \end{pmatrix}$	2.17

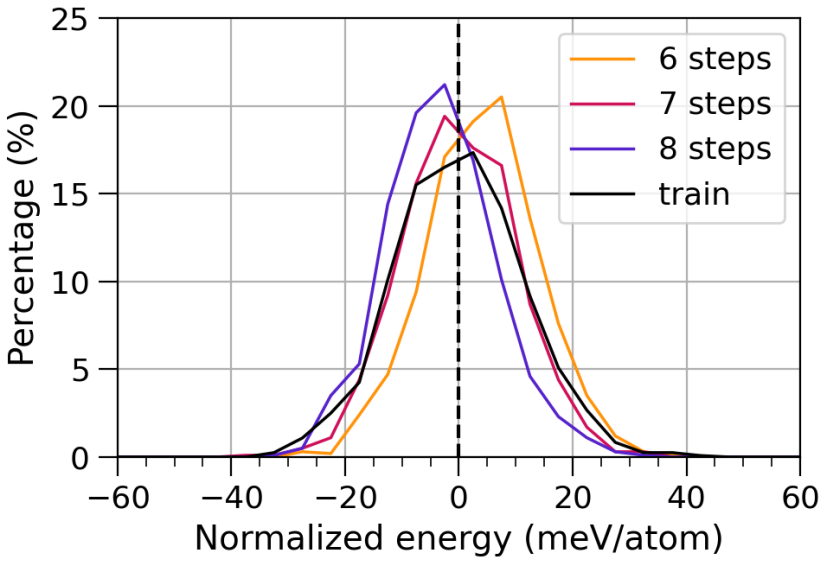

Figure S22: Normalized energy distributions of (relaxed) CDVAE-generated amorphous Si structures, compared to the training structures. The structures were generated using 6, 7, or 8 steps per noise level.

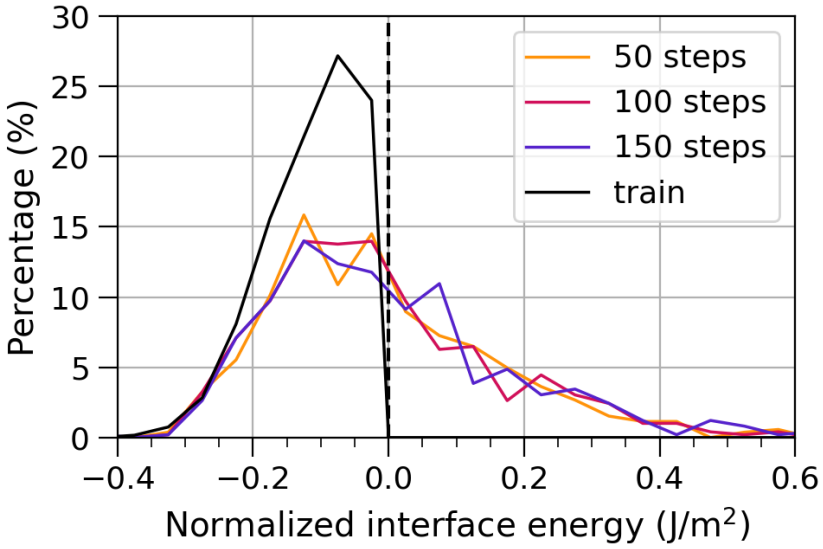

Figure S23: Normalized interface energy distributions of (relaxed) CrysTens-generated interface structures, compared to the training structures. The structures were generated using 50, 100, or 150 time steps.

Table S5: Hyperparameters used for the UniMat 3D U-Net model. The hyperparameter names correspond to those defined in ref. 2. Default values were used for all other hyperparameters.

Hyperparameter	Value
dim	64
dim_mults	(1, 2, 4)
num_resnet_blocks	3
layer_attns	(False, True, True)
layer_cross_attns	(False, True, True)

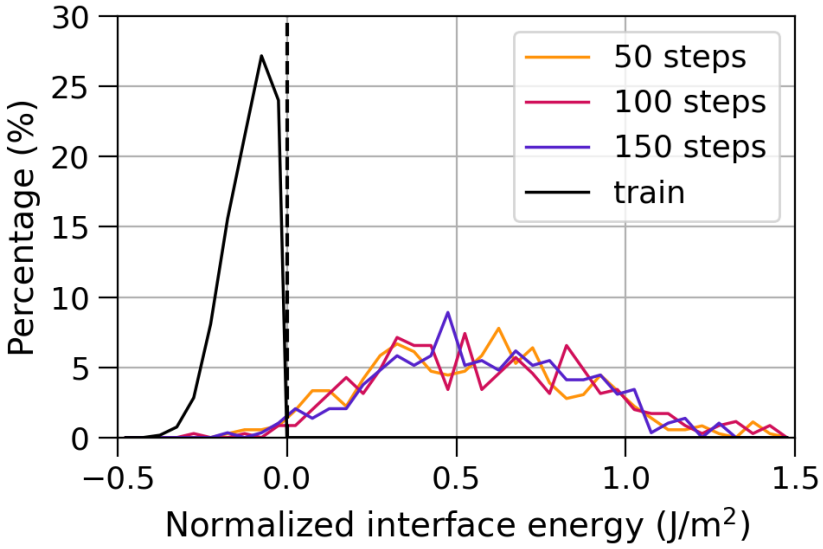

Figure S24: Normalized interface energy distributions of (relaxed) UniMat-generated interface structures, compared to the training structures. The structures were generated using 50, 100, or 150 time steps.

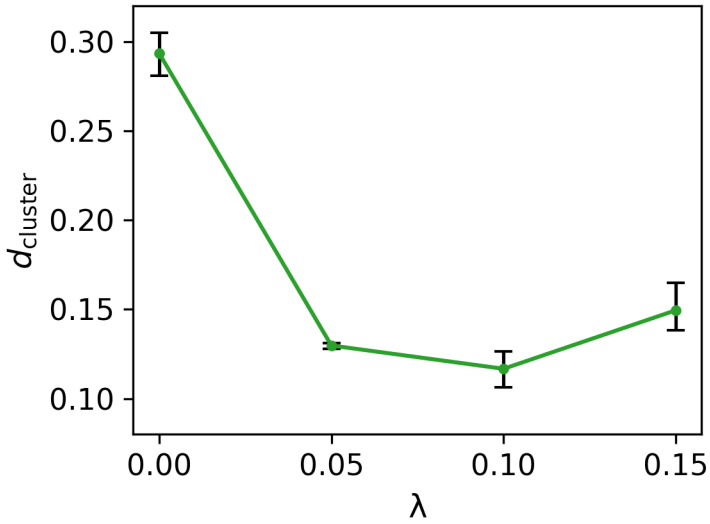

Figure S25: Euclidean distance of the cluster probability fingerprint for alloy structures (300 K, narrow SRO) generated using CryinGAN trained with different λ values.

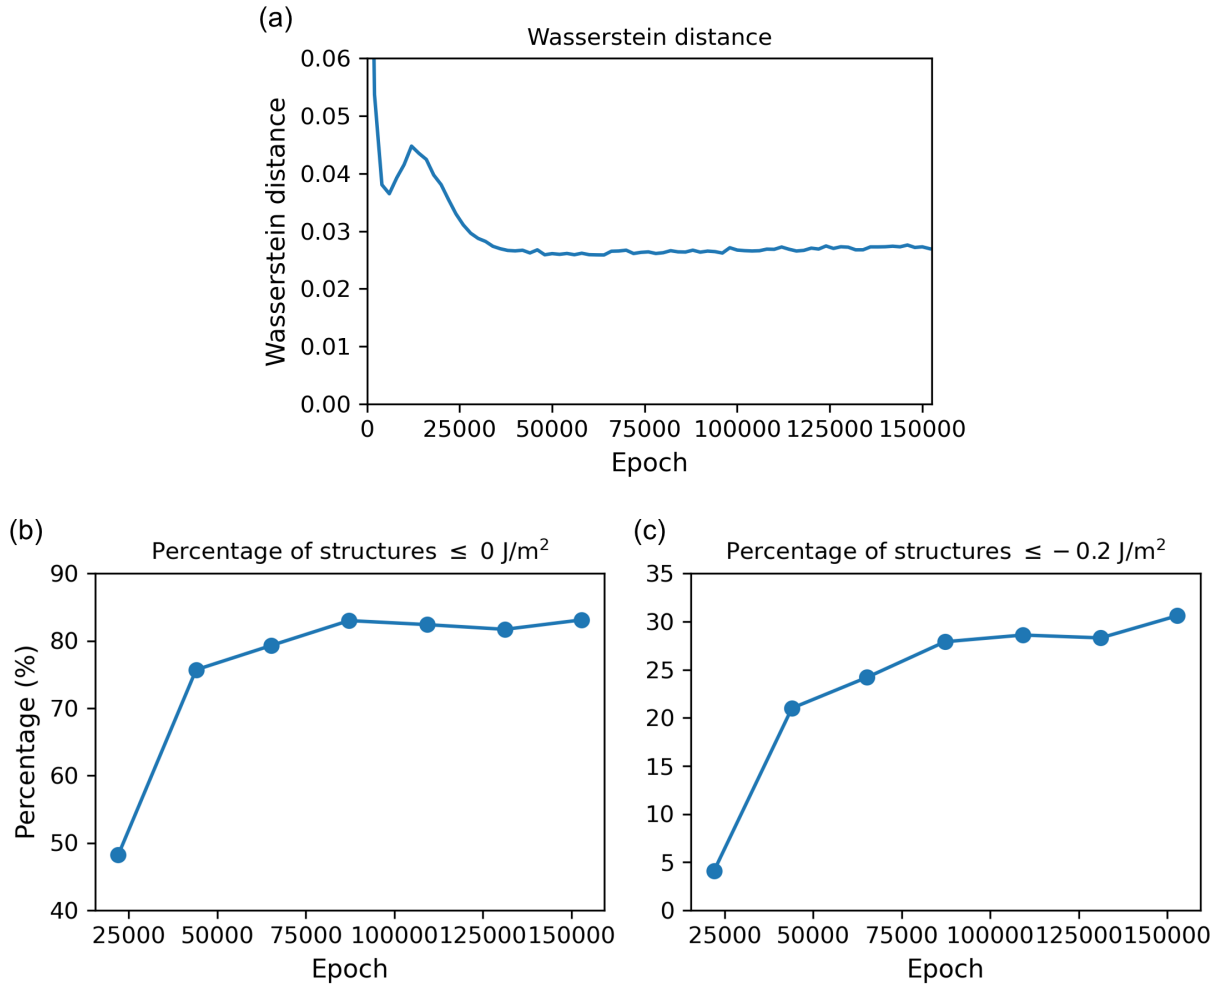

Figure S26: (a) Wasserstein distance as a function of epoch for CryinGAN trained with $\lambda = 0.05$. Note that the Wasserstein distance shown here does not include the gradient penalty term of the Wasserstein loss function. (b) Percentage of relaxed structures with normalized interface energy ≤ 0 J/m^2 as a function of epoch. (c) Percentage of relaxed structures with normalized interface energy ≤ -0.2 J/m^2 as a function of epoch. Structures were generated using a CryinGAN model (trained for a certain number of epochs) and relaxed using the M3GNet interatomic potential. Each data point was calculated using 1,000 relaxed structures.

Table S6: Validation set mean absolute errors (MAEs) and losses of M3GNet models trained with different learning rates and batch sizes. The loss, L, is as defined in the main text. For each model, the epoch with the smallest L is shown. The model with the smallest L is highlighted in bold font. Note that we also trained a model with a learning rate of 0.005 and a batch size of 4, but we found the training to be unstable so the results are omitted here.

Learning rate	Batch size	Energy MAE (meV/atom)	Force MAE (meV/Å)	Stress MAE (GPA)	Loss, L
0.001	**4**	**2.75**	**20.9**	**0.0151**	**0.0251**
0.0005	4	2.82	22.2	0.0179	0.0268
0.001	2	3.45	21.9	0.0203	0.0273
0.001	6	2.88	21.3	0.0156	0.0257

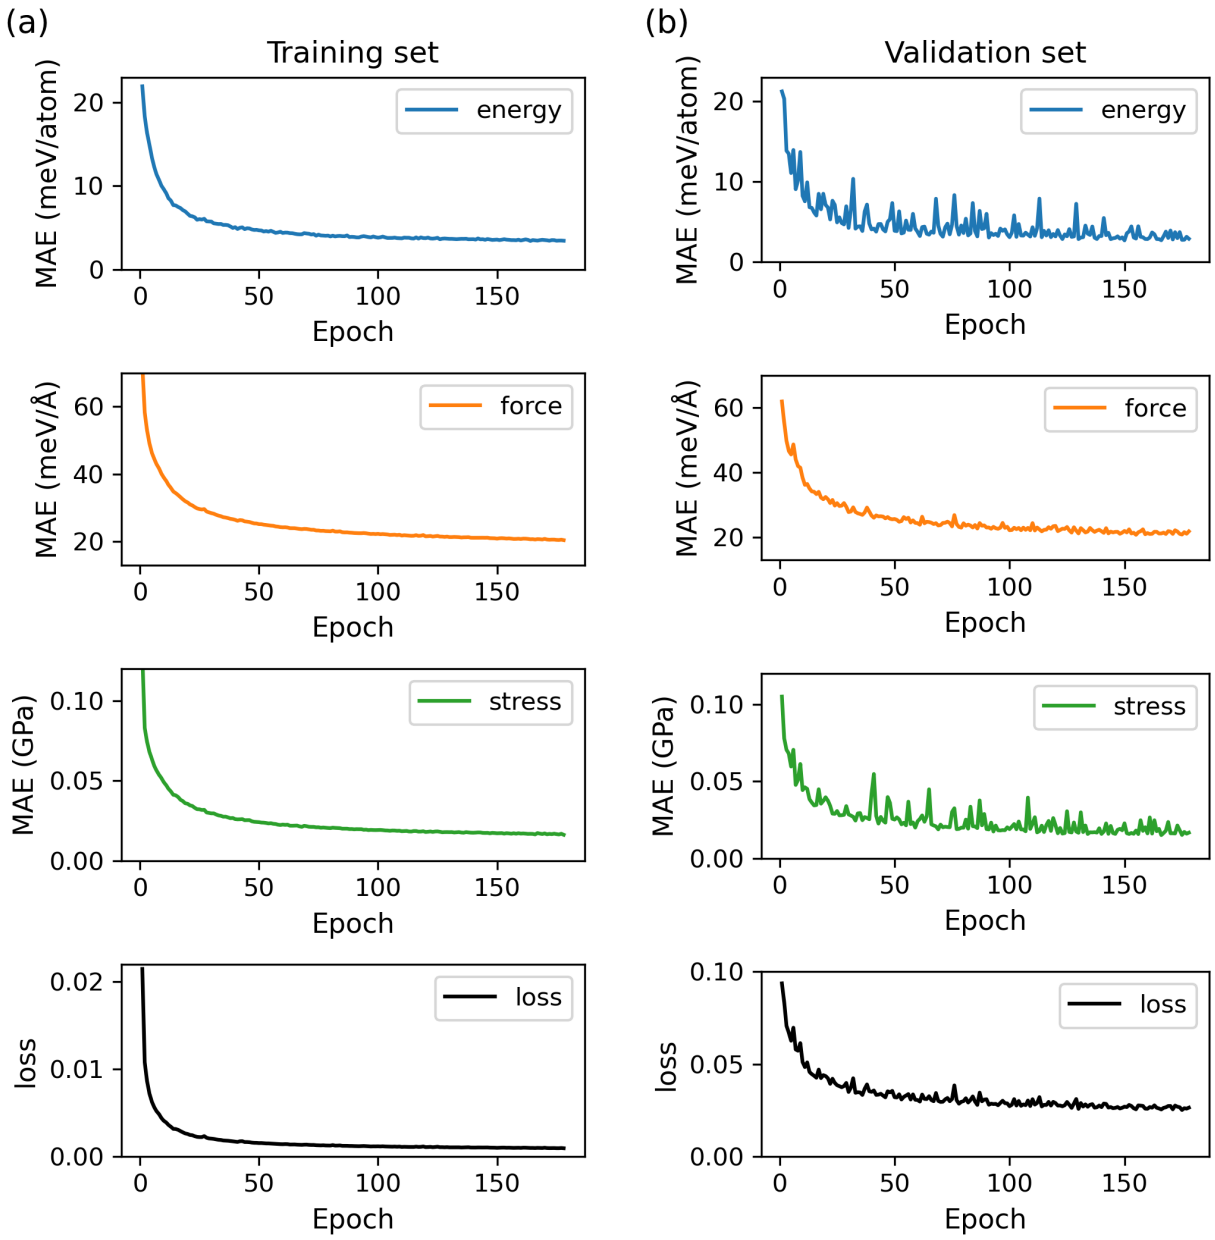

Figure S27: M3GNet training curves of the (a) training set and (b) validation set for the model trained with a learning rate of 0.001 and a batch size of 4. The mean absolute errors (MAEs) for energy, force, and stress, as well as the loss function are plotted against training epoch.

Supplementary Note 1: Oxide-chloride interface disorder mechanism

To investigate the origin of the interfacial disorder observed for oxide-chloride interfaces, interfaces of binary materials (Li_2O-Li, Li_2O-LiCl, and Li_2O-MgS) were studied. All DFT calculations were performed using the same procedure as described in the main text. The Li ($1s^2$ $2s^1$), Cl ($3s^2$ $3p^5$), O ($2s^2$ $2p^4$), Mg ($2p^6$ $3s^2$), and S ($3s^2$ $3p^4$) electrons were treated as valence electrons in the pseudopotentials. Structural relaxations were first performed on unit cells of Li_2O ($Fm\bar{3}m$), Li ($Im\bar{3}m$), LiCl ($Fm\bar{3}m$), and MgS ($Fm\bar{3}m$). The cell shapes, cell volumes, and atom positions were allowed to relax, until the force on each atom was below 0.001 eV/Å. The Brillouin zone was sampled using a (8×8×8) Monkhorst-Pack k-point grid for Li, and a (6×6×6) Monkhorst-Pack k-point grid for Li_2O, LiCl, and MgS. The interface structures were constructed using the MPInterfaces package[3], which implements the lattice matching algorithm proposed by Zur et al.[4]. The configurations of interfaces constructed and their lattice mismatches are listed in Table S4. All interfaces were constructed with vacuum spacings of at least 14 Å. We chose the (100) orientation as a representative plane for Li, LiCl, and MgS.

To study the effect of different terminations of Li_2O surfaces on the interfacial structure with LiCl, we chose the Li_2O(100) orientation which can either be Li-terminated of O-terminated. The surfaces of the Li_2O(100) slabs are polar, so half of the Li/O atoms were moved from one surface to the other to neutralize the polarity (resulting in 'Tasker Type 2b' surfaces[5,6]). Structural relaxations were performed on Li_2O(100)-LiCl(100) interfaces for both terminations, allowing the cell shapes, cell volumes, and atom positions to relax until the force on each atom was below 0.05 eV/Å. The Brillouin zone was sampled using a (3×2×1) gamma-centered k-point grid. To study the effect of mechanical stiffness on interfacial structure, we chose the Li_2O(110) orientation, which exposes both Li and O at its surface, and paired it with Li(100), LiCl(100), and MgS(100). Structural relaxations were performed on these interfaces using the same procedure. The Brillouin zone was sampled using a (8×2×1), (3×2×1), and (3×2×1) gamma-centered k-point grid for Li_2O(110)-Li(100), Li_2O(110)-LiCl(100), and Li_2O(110)-MgS(100) respectively.

While many heterointerfaces can adopt regular epitaxial registries connecting two materials with well-defined crystalline orientations, other heterointerfaces show irregular, disordered interfacial patterns. We found that chlorides have an innate tendency to form disordered interfaces with oxides. As shown in Fig. S28, using binary instead of ternary materials, disordered interfacial structures are obtained with only the combination of Li_2O and LiCl. We suggest two interdependent reasons for their occurrence. The first reason is the bond formation between O (in the oxide) and Li (in the chloride). Fig. S28a shows the interface structure of Li_2O(100)-LiCl(100) for Li_2O(100) slabs that are either Li-terminated or O-terminated. When Li_2O(100) is Li-terminated, there is minimal rearrangement at the interface with only some Li-Cl bond formation. In contrast, when Li_2O(100) is O-terminated, significant atomic rearrangement is observed with a combination of Li-O and Li-Cl bond formation. The stronger interaction between O and Li is likely due to the higher charge density of O compared to Cl. However, Li-O bond formation alone does not lead to highly disordered interfaces, since such degree of disorder is not observed in other O-containing interfaces such as oxide-oxide and oxide-sulfide interfaces[7,8].

The second reason that the chlorides can form disordered interfaces is due to their intermediate mechanical stiffness. To form bonds with the O in the oxide slab, the other slab must distort its crystal structure to align its atoms with O, and the stiffness of the material determines the extent of the distortion allowed. Fig. S28b shows the interface structure between Li_2O and materials in increasing order of mechanical stiffness (Li, LiCl, and MgS). MgS has the same rock salt crystal structure as LiCl but with higher stiffness[9,10]. Due to its higher stiffness, it is unable to distort significantly to bond with O. On the other hand, LiCl is soft enough that it can distort itself, whilst maintaining some of its structure, resulting in a more distorted interface structure. For a very soft material like Li metal[11], the atoms can rearrange themselves relatively freely to bond with O, so large voids such as those observed for LiCl are less likely to form. The intermediate stiffness of the chlorides, combined with bond formation with O, results in the chloride solid electrolytes' ability

to form highly disordered interface structures with the oxide cathode. The disordered nature of the $Li_3ScCl_6(100)$-$LiCoO_2(110)$ interfaces is therefore understood from the two considerations above.

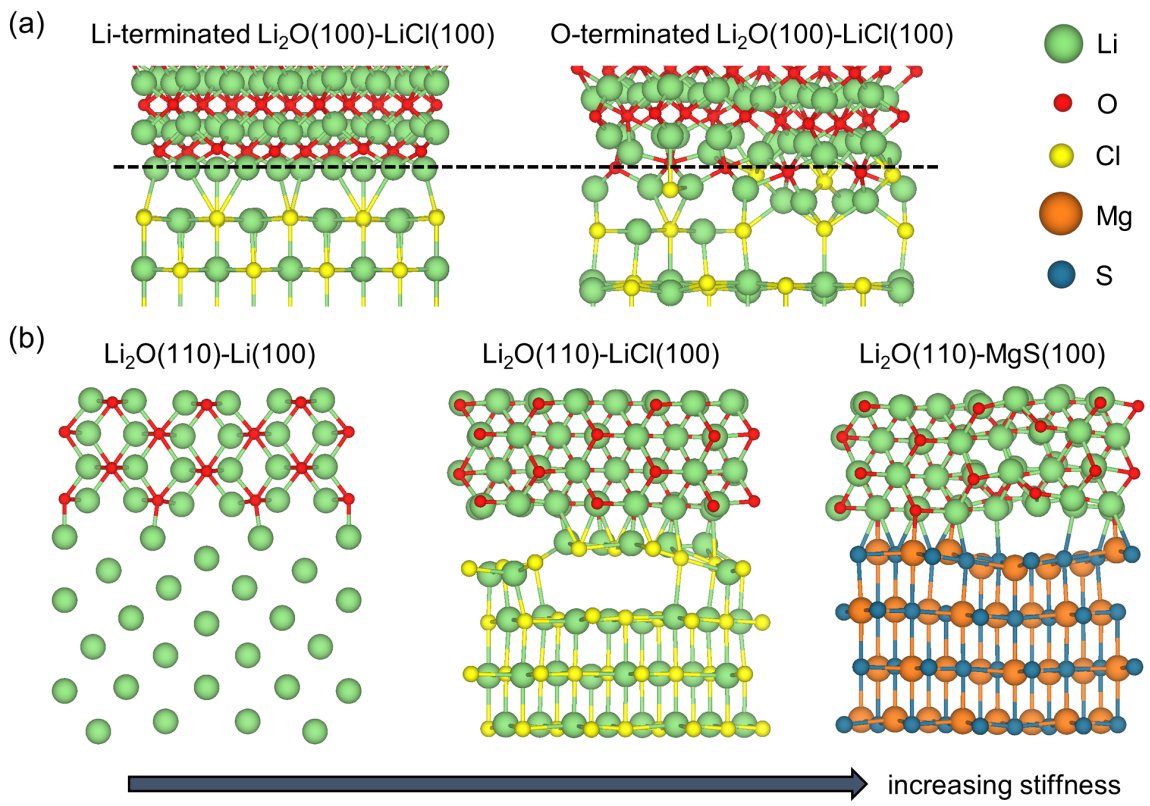

Figure S28: (a) Interface structures of $Li_2O(100)$-$LiCl(100)$ for Li-terminated and O-terminated $Li_2O(100)$. The dashed line indicates the termination layer of $Li_2O(100)$. (b) Interface structures for $Li_2O(110)$-$Li(100)$, $Li_2O(110)$-$LiCl(100)$, and $Li_2O(110)$-$MgS(100)$. The mechanical stiffness of the bottom slab material increases from left to right.

Supplementary Note 2: Alternative CryinGAN architectures

We considered a different architecture that circumvents the need to tune λ in CryinGAN, where both fractional coordinate and bond distance discriminators were combined into a single discriminator, which we refer to as CryinGAN-comb (see Fig. S29). The outputs were combined after the pooling layer, and the model was allowed to learn the relative importance of the coordinate and bond distance latent features on its own through fully connected layers. For CryinGAN-comb, the bond distances were directly obtained using fractional coordinates (instead of Cartesian coordinates), to provide a more direct link to the atomic coordinates which are also represented in fractional coordinates. Structures were generated using CryinGAN-comb models and relaxed using M3GNet. Compared to CryinGAN trained with $\lambda = 0$, CryinGAN-comb generates structures with fewer atoms too close together (around 2 times fewer pairs of atoms $\leq$ 1.5 Å). However, as shown in Fig. S30, the interface energy distribution of (relaxed) structures generated from CryinGAN-comb is significantly higher than CryinGAN. These results show that combining the discriminators in CryinGAN-comb leads to less useful discriminator gradients to train the generator, resulting in poor quality of structures being generated. In contrast, separating the discriminators as in CryinGAN helps the GAN to learn more effectively.

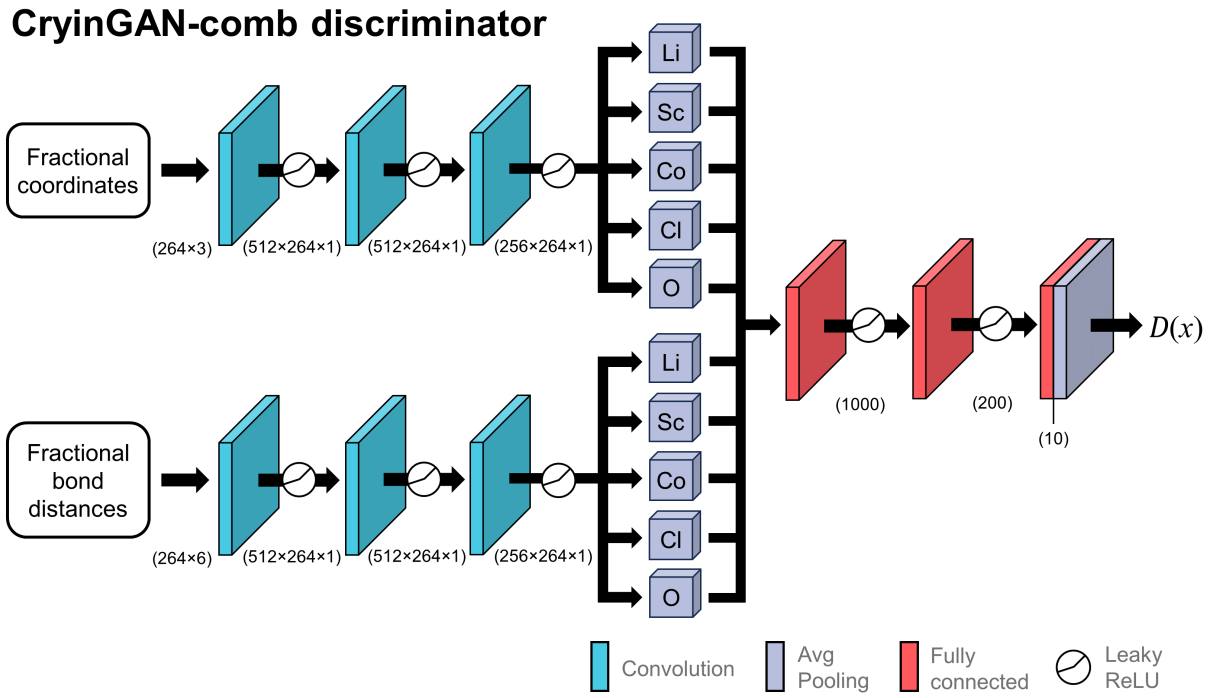

Figure S29: Discriminator architecture of CryinGAN-comb. The output of the two discriminators of CryinGAN are combined after pooling.

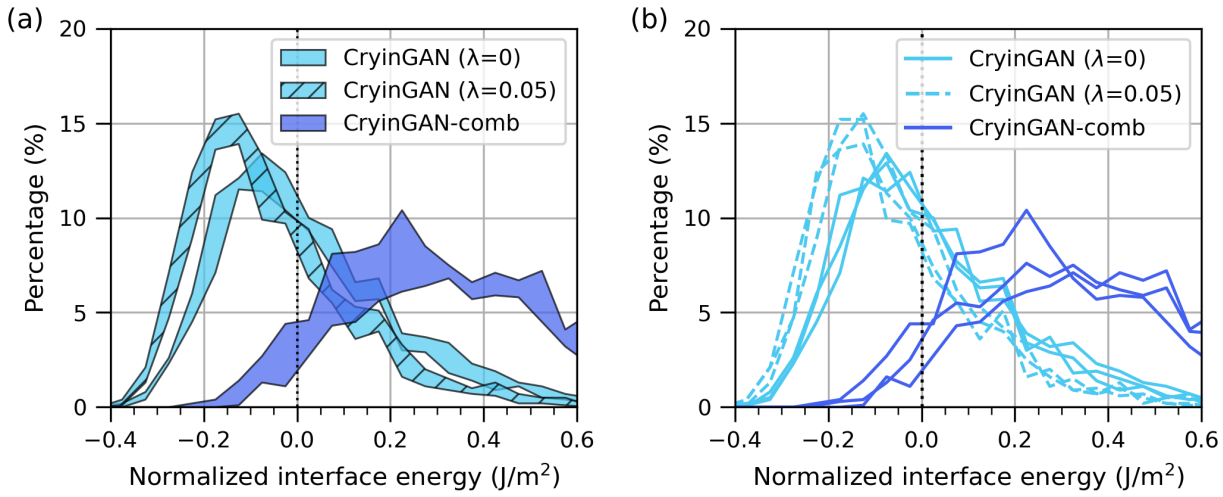

Figure S30: Normalized interface energy distributions of structures generated using CryinGAN (λ = 0, 0.05) and CryinGAN-comb. For each model configuration, 3 separate models were trained. The energy distribution are shown as shaded in (a) and unshaded in (b). All structures were relaxed using M3GNet, and the interface energies shown are based on M3GNet-calculated energies.

Next, we considered different choices of the pooling layer for the discriminator. Wang et al.[12] found that the type of pooling operation affected the sampling sensitivity of the discriminator and the overall performance of the GAN. The sampling sensitivity describes how sensitive the discriminator is to changes in point density or the sampling pattern of the input point cloud. Their results suggest that max pooling produces a discriminator with lower sampling sensitivity than average pooling. The CryinGAN discriminators use average pooling, and we tested an architecture which uses max pooling instead, referred to as CryinGAN-max (see Fig. S31a). We also tested an architecture named CryinGAN-mix that uses the mix pooling operation proposed by Wang et al.[12], where both max and average pooling operations are used together (see Fig. S31b). CryinGAN-max and CryinGAN-mix were trained using $\lambda = 0$ and 0.05. For $\lambda = 0.05$, only the pooling layer of the bond distance discriminator was changed, and we kept the average pooling for the fractional coordinate discriminator. This choice allows us to study the effect of the pooling choice on the two discriminators independently.

The interface energy distributions of relaxed structures generated using CryinGAN, CryinGAN-max, and CryinGAN-mix are shown in Fig. S32. For $\lambda = 0$ (Fig. S32a-b), we observe that CryinGAN (average pooling) significantly outperforms CryinGAN-max and CryinGAN-mix. Although the use of max pooling was beneficial for the generation of 3D objects where the object shape is the most important aspect to capture[12], it appears that a higher sampling sensitivity is needed for atomic configurations, for which the local coordination environment is the important aspect to capture. For $\lambda = 0.05$ (Fig. S32c-d), we observe that the distributions are more similar across the three architectures, but CryinGAN still has the lowest energy distribution. The pooling choice of the bond distance discriminator does not relate to the sampling sensitivity like the fractional coordinate discriminator, and the choice appears to have a smaller effect on model performance. Overall, we find that average pooling works best for both discriminators.

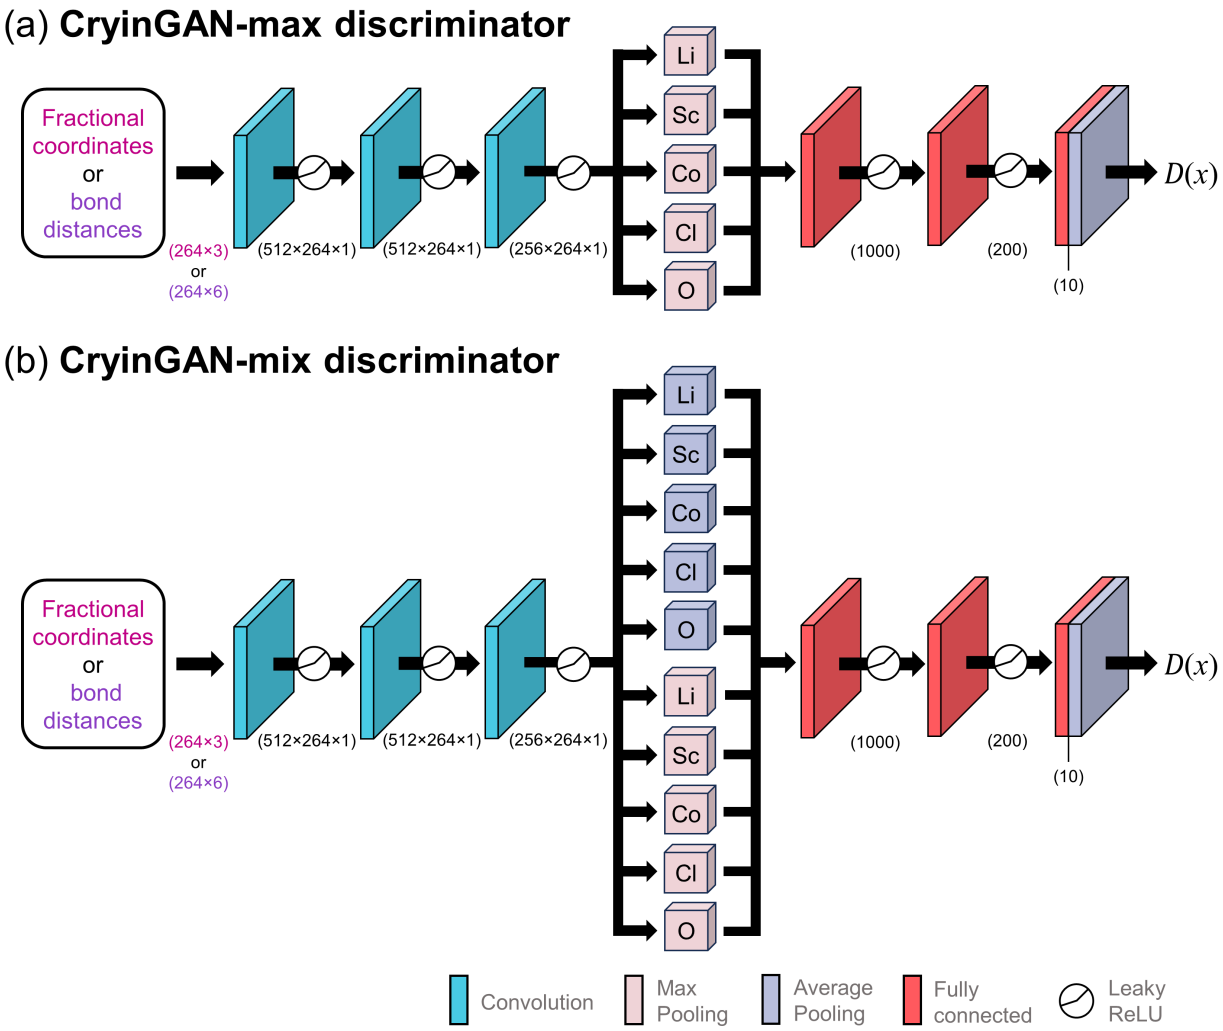

Figure S31: Discriminator architectures of (a) CryinGAN-max and (b) CryinGAN-mix. When the bond distance discriminator is not used, max/mix pooling is applied to the fractional coordinate discriminator. When the bond distance discriminator is used, max/mix pooling is applied to the bond distance discriminator, and average pooling is applied to the fractional coordinate discriminator.

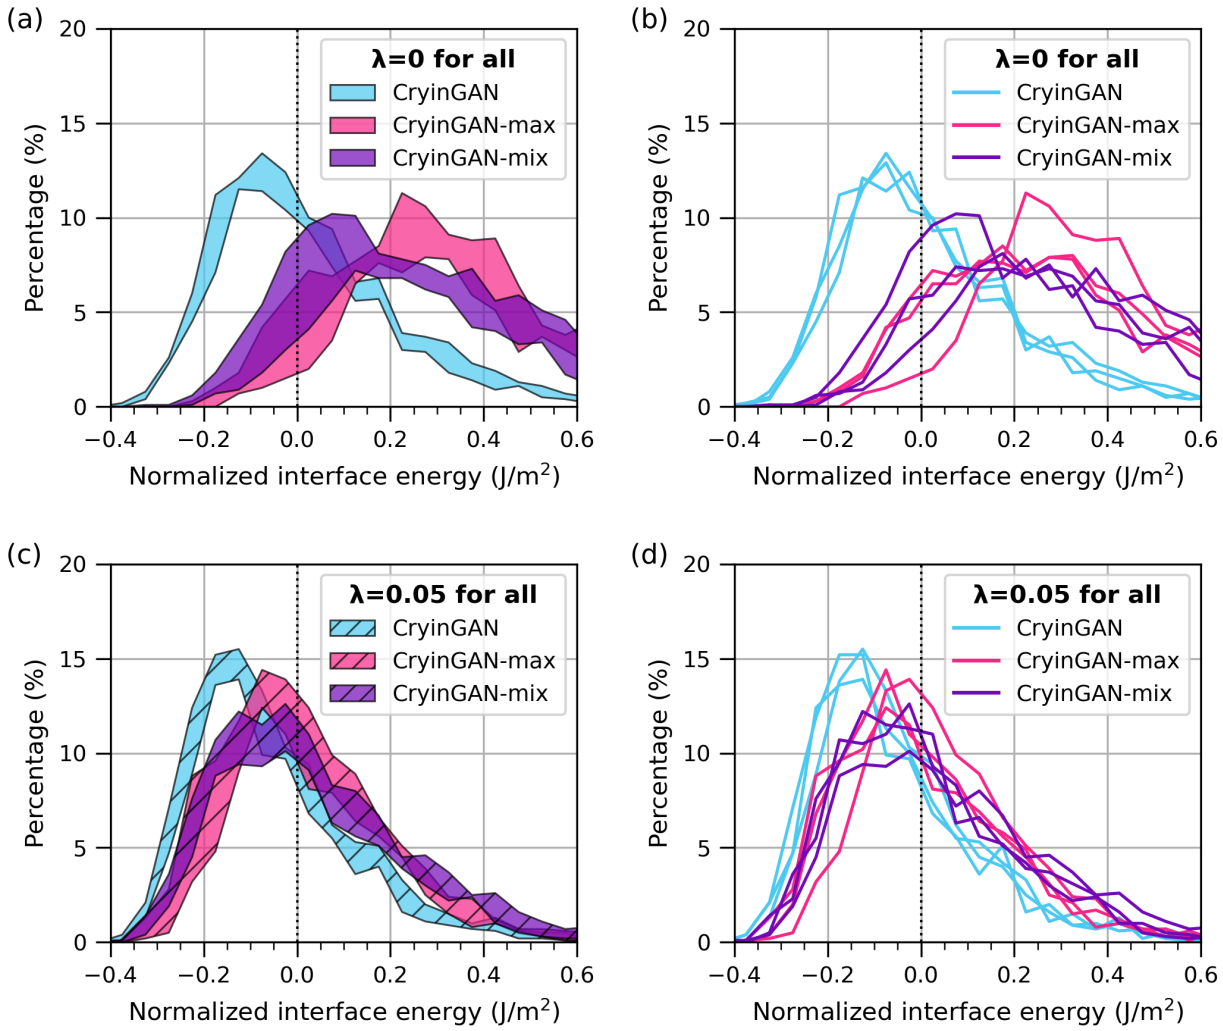

Figure S32: Normalized interface energy distributions of structures generated using CryinGAN (average pooling), CryinGAN-max (max pooling), and CryinGAN-mix (mix pooling). (a) and (b) show models without the bond distance discriminator ($\lambda = 0$), where pooling is varied for the fractional coordinate discriminator. (c) and (d) show models with the bond distance discriminator ($\lambda = 0.05$), where pooling is varied for the bond distance discriminator only (average pooling is used for the fractional coordinate discriminator). For each model configuration, 3 separate models were trained. The energy distribution are shown as shaded on the left, and unshaded on the right. All structures were relaxed using M3GNet, and the interface energies shown are based on M3GNet-calculated energies.

Supplementary Note 3: Interface Li coordination motif analysis of CryinGAN-generated structures

Table S7 shows the Euclidean distance and cosine similarity between the average interface Li fingerprint of the CryinGAN/high energy dataset and the training dataset. Compared to the high energy dataset, the CryinGAN dataset has a slightly lower Euclidean distance to the low energy dataset, and similar cosine similarity. This small difference is better understood by examining the distribution of the most likely Li coordination motif as shown in Fig. S33. Compared to the coordination motif distribution of the training dataset, the distribution of the high energy dataset is shifted upwards towards motifs with lower coordination number, indicating that fewer bonds are leading to higher energies. On the other hand, the coordination motif distribution of the CryinGAN-generated structures shows a higher similarity to the training structures, reflecting the smaller motif fingerprint Euclidean distance seen in Table S7.

Table S7: Euclidean distance and cosine similarity between the average interface Li site fingerprint of the training structures and the CryinGAN/high-interface-energy structures. The 95 % bootstrap confidence intervals are shown in brackets.

Dataset	Euclidean distance (95 % CI)	Cosine similarity (95 % CI)
CryinGAN	0.1297 (0.1106 to 0.1455)	0.9974 (0.9968 to 0.9982)
High energy	0.1361 (0.1226 to 0.1471)	0.9973 (0.9969 to 0.9979)

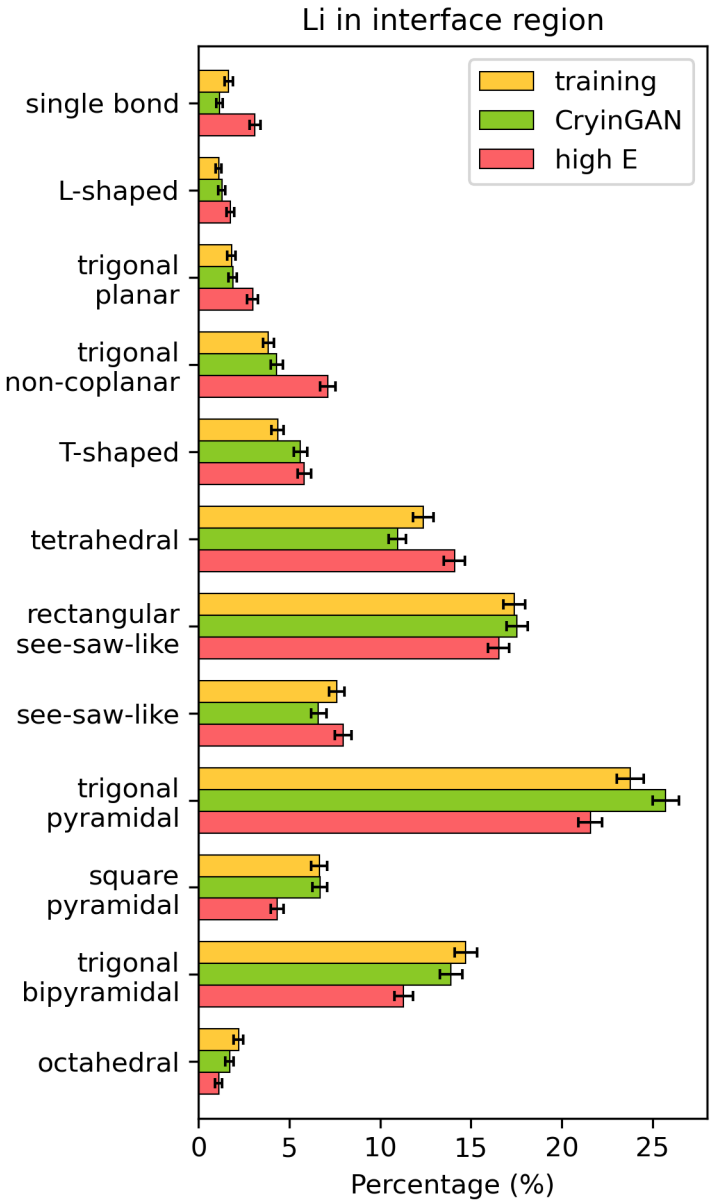

Figure S33: Coordination motif distributions of Li in the interface region for three datasets: (1) training structures with low interface energy, (2) CryinGAN-generated structures, and (3) structures with high interface energy. All structures were relaxed using M3GNet followed by DFT calculations. Error bars represent 95 % bootstrap confidence intervals. The coordination motifs are ordered in increasing coordination number from top to bottom.

Supplementary Note 4: RDF analysis of CryinGAN-generated interfaces
The Sc-O, Sc-Sc, and Li-Li RDFs are shown in Fig. S34. For the Sc-O RDF (Fig. S34a), the high energy dataset exhibits peaks with lower magnitude, reflecting its reduced degree of Sc-O bonding. On the other hand, the Sc-Sc RDF (Fig. S34b) shows a higher peak around 4 Å, indicating that the high energy structures have a higher proportion of Sc cations closer together. The associated Sc-Sc repulsion raises the energy of the structures. In the Li-Li RDF (Fig. S34c), a new peak appears around 2.5 Å. This short Li-Li interatomic distance again raises the energy of the structures due to Li-Li repulsion. In contrast, the CryinGAN dataset does not show any of these high-energy features, and its RDFs are similar to the training dataset.

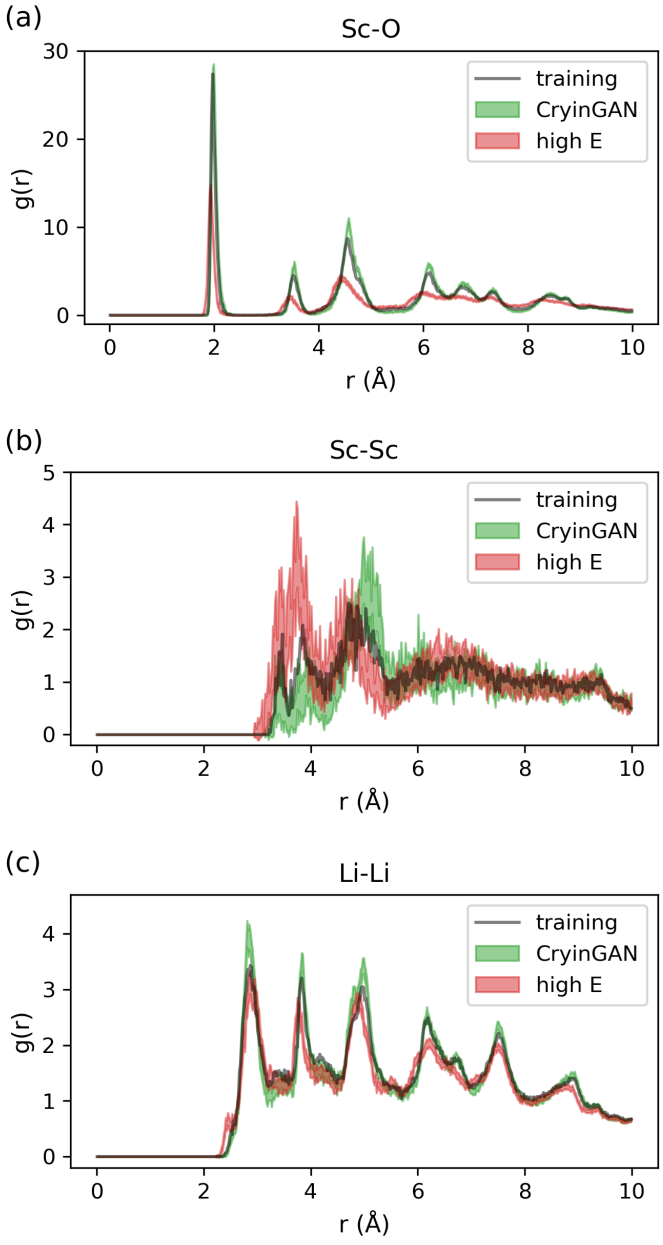

Figure S34: Radial distribution functions of atoms in the interface region for (a) Sc-O, (b) Sc-Sc, and (c) Li-Li. The thickness of each curve represents the 95 % bootstrap confidence interval.

References

[1] T. Su, B. J. Blankenau, N. Kim, J. A. Krogstad and E. Ertekin, Acta Materialia, 2024, **276**, 120088.

[2] P. Wang, imagen-pytorch, 2024, `https://github.com/lucidrains/imagen-pytorch`.

[3] K. Mathew, A. K. Singh, J. J. Gabriel, K. Choudhary, S. B. Sinnott, A. V. Davydov, F. Tavazza and R. G. Hennig, Computational Materials Science, 2016, **122**, 183–190.

[4] A. Zur and T. C. McGill, Journal of Applied Physics, 1984, **55**, 378–386.

[5] P. W. Tasker, Journal of Physics C: Solid State Physics, 1979, **12**, 4977–4984.

[6] J. D. Gale and A. L. Rohl, Molecular Simulation, 2003, **29**, 291–341.

[7] B. Gao, R. Jalem and Y. Tateyama, ACS Applied Materials & Interfaces, 2021, **13**, 11765–11773.

[8] S. He, K. Chen, M. Saunders, J. Li, C. Q. Cui and S. P. Jiang, Journal of The Electrochemical Society, 2017, **164**, F1437.

[9] J. Wang, M. Deng, Y. Chen, X. Liu, W. Ke, D. Li, W. Dai and K. He, Materials Chemistry and Physics, 2020, **244**, 122733.

[10] Y.-D. Guo, Z.-J. Yang, Q.-H. Gao and W. Dai, Physica B: Condensed Matter, 2008, **403**, 2367–2371.

[11] A. Masias, N. Felten, R. Garcia-Mendez, J. Wolfenstine and J. Sakamoto, Journal of Materials Science, 2019, **54**, 2585–2600.

[12] H. Wang, Z. Jiang, L. Yi, K. Mo, H. Su and L. J. Guibas, Rethinking Sampling in 3D Point Cloud Generative Adversarial Networks, 2020.