

Cite this: *Digital Discovery*, 2026, 5, 1590

The loss landscape of powder X-ray diffraction-based structure optimization is too rough for gradient descent

Nofit Segal, ^a Akshay Subramanian, ^a Mingda Li, ^b Benjamin Kurt Miller ^c and Rafael Gómez-Bombarelli ^{*a}

Solving crystal structures from powder X-ray diffraction (XRD) is a critical inverse problem in materials characterization. This work studies the mapping from diffractogram to crystal structure using gradient-based optimization of the structure with XRD similarity as the objective, and evaluates the retrieval of ground-truth geometries from moderately distorted structures. We find that commonly used XRD similarity metrics result in an ill-posed, highly non-convex loss landscape where high signal agreement does not necessarily imply structural accuracy, a phenomenon driven by spurious peak overlaps. Constraining the optimization to the ground-truth crystal family significantly improves retrieval, and yields higher correlation between structural similarity and XRD similarity. Nevertheless, the landscape may remain non-convex along certain symmetry axes. Finally, we contrast this with the interatomic potential energy landscape, which exhibits smooth, locally convex behavior for identical structural perturbations.

Received 15th January 2026
Accepted 31st March 2026

DOI: 10.1039/d6dd00017g

rsc.li/digitaldiscovery

1 Introduction

Determining the atomic structure of a crystal from its powder X-ray diffraction (XRD) pattern is a longstanding and central challenge in materials characterization.^{1,2} The inverse problem of recovering the full three-dimensional crystal structure solely from an XRD pattern is extremely challenging due to the loss of phase information of the scattered waves—known as the phase problem.^{3,4} Nevertheless, powder diffraction is widely used for identifying and characterizing crystalline solids. In practice, this is typically achieved by comparing the observed XRD spectrum to a reference database and performing least-squares refinement, known as Rietveld analysis.^{5–7} However, Rietveld analysis is highly sensitive to initial parameters,⁸ and more importantly, it relies on the presence of the correct structure in the database and cannot be used to reconstruct novel or unreported phases.

Experimental phenomena such as preferred orientation, peak overlap, crystal twinning, and instrumental noise further complicate structural determination from powder XRD patterns.^{9–11} Furthermore, the formation of phase mixtures further complicates interpretation, as diffraction peaks from multiple phases become superimposed.^{4,12}

From a computational perspective, structural ambiguity remains even under idealized experimental conditions. Distinct compositions can exhibit highly correlated XRD patterns when they share similar crystal symmetry and lattice parameters, a behavior that is common in both minerals and metallic alloys. This frequently occurs in solid solutions, where the same atomic species occupy crystallographic sites but with slightly different stoichiometric ratios, leading to only minor lattice perturbations and nearly indistinguishable diffraction patterns across a composition range.^{4,13–18} A related but distinct source of ambiguity arises in isomorphous compounds, in which one atomic species is fully substituted by another—typically of similar valence and ionic radius—while preserving the same crystal structure and lattice geometry.^{17,18} In both cases, different compositions map to highly similar diffraction signatures, rendering the XRD-to-structure mapping inherently one-to-many.

Moreover, even when the atomic composition is fixed, a material may crystallize into different phases with related but distinct symmetries. Such polymorphs can be described by different space groups yet still produce strongly correlated diffraction patterns, further complicating structural inference.^{2,9}

Notably, small distortions in lattice parameters or displacements of atomic coordinates do not induce uniformly smooth changes in the diffraction signal. While some distortions lead to gradual shifts in peak positions or intensities, others can trigger discontinuous changes, such as peak splitting, appearance, or disappearance, through abrupt changes in Bragg conditions.⁴

^aDepartment of Materials Science and Engineering, MIT, Cambridge, MA, USA. E-mail: rafagb@mit.edu^bDepartment of Nuclear Science and Engineering, MIT, Cambridge, MA, USA^cFAIR, Meta, San Francisco, USA

This introduces a *non-smooth* relationship between structure and XRD signal.

Recently, there has been a surge of interest in crystal structure determination from XRD patterns using generative modeling.^{19–25} A growing body of work applies gradient-based optimization approaches that leverage differentiable physics to refine generated or otherwise-obtained crystal structures by minimizing the difference between simulated and target XRD patterns. For example, Riesel *et al.*²⁰ generate crystals conditioned on a given XRD pattern and post-process them using a differentiable XRD simulator to update lattice parameters *via* gradient descent (GD). Parackal *et al.*²⁶ systematically enumerate candidate crystals given composition and space group inputs, and restricts the GD optimization to atomic positions along Wyckoff degrees of freedom. Lee *et al.*²⁷ create candidate crystals using an evolutionary algorithm, followed by crystals morphing by maximizing the cosine similarity between the XRD patterns. Outside the powder diffraction setting, GD has further been applied to determine lattice parameters from single-crystal diffraction patterns.²⁸

As gradient-based refinement relies on comparing simulated and target diffraction patterns, recent work has also focused on developing more robust XRD-similarity metrics. Otero-de-la Roza²⁹ introduced a cross-correlation-based metric that captures equivalence between diffraction patterns while remaining invariant to lattice distortions. Building on this work, Racioppi *et al.*³⁰ applied the metric to crystal structure prediction from XRD data, jointly optimizing the structure by minimizing both this similarity metric and the structure's enthalpy. Hernández-Rivera *et al.*³¹ systematically analyzed the sensitivity of different families of similarity metrics under isotropic lattice strain. Li *et al.*³² proposed an entropy-based similarity measure for spectra and demonstrated its utility for molecular database retrieval from mass spectrometry data.

In this work, we explore the powder XRD-to-structure mapping through the lens of GD optimization. The task is to recover correct structures from moderately deviated ones based solely on XRD similarity. This work investigates whether the XRD landscape is locally smooth enough for GD to guide back to the correct configuration.

Inspired by experimentally observed symmetry-breaking effects such as thermal expansion from lattice vibrations and thermal fluctuations, we introduce two types of distortions: random lattice distortions and uncorrelated atomic displacements.^{8,33–35} These distortions resemble crystal structures predicted by generative models, which often produce nearly correct geometries but with imperfect symmetry.^{36–38} Through this study, we examine the challenge of “the last mile” in structure elucidation from XRD.

We find that mapping XRD patterns to crystal structures is challenging because high diffraction agreement, as currently measured in literature, does not ensure geometrical accuracy. We show that commonly used XRD similarity measures, such as cosine similarity, mean squared error (MSE), entropy similarity, and cross-correlation, are sensitive to both lattice and coordinate noise distortions, and optimization between distorted structures and ground-truth XRD diffraction can become trapped in local minima.

2 Results and discussion

Optimization was performed on distorted crystal structures across a range of noise types and levels. For each condition, 50 distorted versions were generated for each of 10 ground truth structures, yielding 500 distorted inputs per noise setting.

Fig. 1 illustrates match rates obtained through crystal structure optimization w.r.t XRD similarity metrics. For lattice distortions, the largest drop in match rate occurs between noise levels of 0.05 and 0.1. Although 0.1 corresponds to the lattice tolerance used by the structure matching criterion (see Section 4), the noise level controls the maximum applied strain, such that lattice parameters typically remain within the matching tolerance. Using different similarity objectives makes little to no difference in performance.

2.1 Symmetry constraints: strengths and limitations

Incorporating symmetry-based constraints during XRD-based optimization notably improves robustness to lattice noise for many structures in this study, as shown in Fig. 1a by the higher match rates achieved when constraints are applied. The constraints (see Subsection 4.4) project updates back into the correct crystal family at each optimization step, thereby guiding the search along a reduced-dimensionality symmetry-consistent

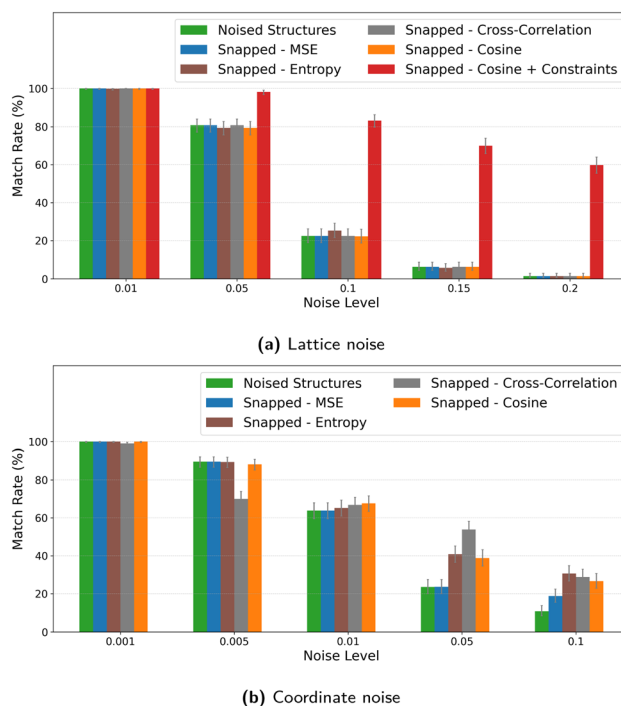


Fig. 1 Results of XRD-based optimization under two types of noise. Crystal structures were optimized with respect to XRD similarity metrics using the snap method,²⁰ which struggles to recover the correct structure under both lattice and coordinate perturbations. The plots show match rates computed with StructureMatcher (ltol = 0.1, stol = 0.2, angle_tol = 5°) under random lattice (a) and coordinate (b) perturbations. Error bars represent 95% Jeffreys binomial credible intervals.³⁹ For lattice distortions, incorporating lattice constraints significantly improves robustness, even at high noise levels.



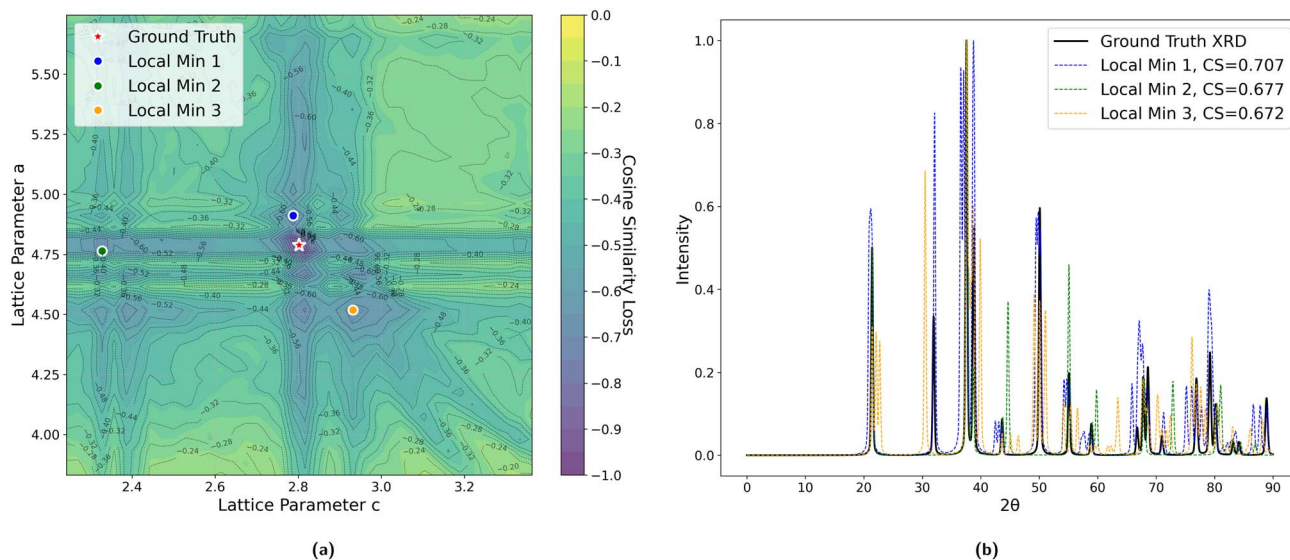


Fig. 2 2D landscape of XRD cosine similarity (CS) loss as a function of lattice parameters a and c [Å] of U_2Ti structure, illustrating the presence of multiple local minima. (a) Cosine similarity loss topographic map showing non-convex behavior with several local minima. (b) XRD patterns for the structures corresponding to the marked local minima: all exhibit reasonably high cosine similarity to the ground truth pattern despite having different lattice parameters.

path and helping the optimizer avoid local minima unrelated to the desired symmetry.

2.1.1 Roughness of simplified loss landscape cross sections. Fig. 2, 3, and S2 illustrate the non-convex nature of the XRD-based loss landscape with respect to the lattice parameters, shown through 2D cross-sections of the optimization surface. These plots represent simplified views of the underlying optimization landscape, which, in the case of a distorted lattice, occurs in a six-dimensional space corresponding to the six lattice parameters.

Fig. 2 shows distortions of the lattice parameters a and c of U_2Ti , a hexagonal structure with space group $P6/mmm$ (no. 191). The heatmap color indicates the cosine similarity between the XRD spectra of the distorted structures and the ground truth. The resulting contour map reveals multiple deep local minima, indicating the optimizer's potential to get trapped in suboptimal solutions. The three most prominent local minima are highlighted, and their corresponding XRD patterns are shown in the right panel. Despite their structural deviation from the true lattice parameters, the patterns show high cosine

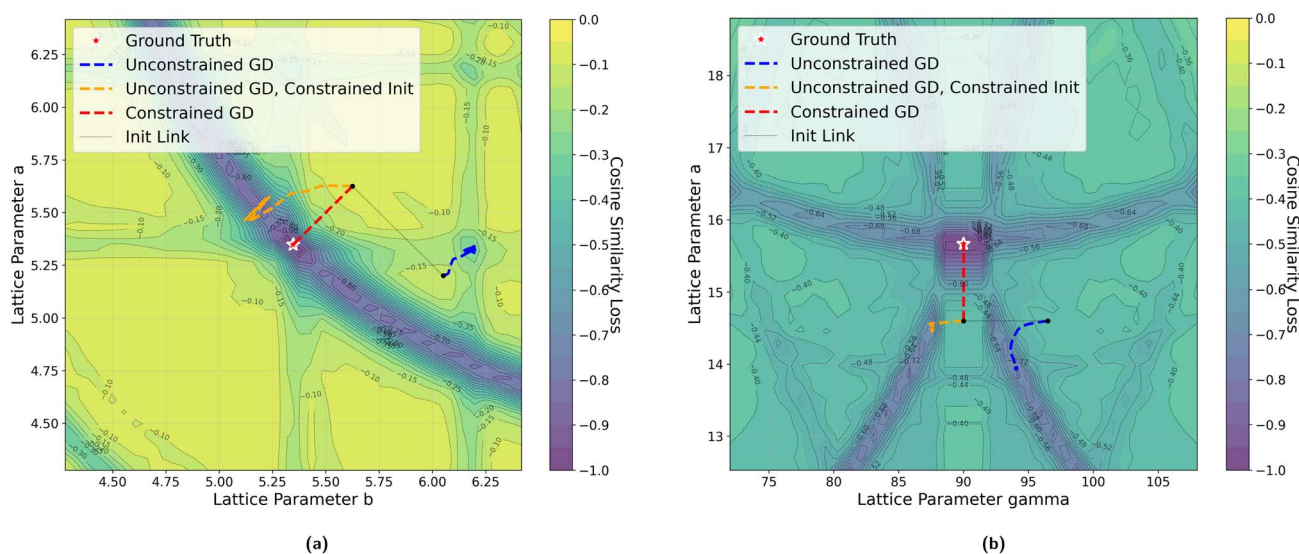


Fig. 3 2D landscape of XRD cosine similarity (CS) loss as a function of lattice parameters, with simulated optimization paths for XRD-based gradient descent (GD). Unconstrained GD, unconstrained GD with a constrained initialization, and fully constrained GD. Unconstrained GD converges to some local minima, even with constrained initialization, whereas constrained GD reaches the ground truth. (a) Lattice parameters a and b [Å] of cubic Au_2S are perturbed. (b) Lattice parameters a [Å] and γ [°] of monoclinic $Na_3MnCoNiO_6$ are perturbed.



similarity to the ground truth due to subtle shifts and peak splittings that preserve the overall spectral profile.

Fig. 3 provides two illustrative examples demonstrating how symmetry constraints can facilitate correct structure reconstruction. In these cases, only a two-dimensional slice of the optimization landscape is visualized for clarity. This is a simplification of the full optimization process, which, for the distorted lattice case, occurs in a six-dimensional space. We show simulated GD trajectories for two representative structures under three settings: (i) unconstrained GD, (ii) unconstrained GD initialized at a constrained point, and (iii) fully constrained GD.

In Fig. 3a, we perturb the lattice parameters a and α of Au_2S , a cubic structure with space group $Pn\bar{3}m$ (no. 224). The unconstrained GD trajectory converges to a distant local minimum, whereas unconstrained GD with constrained initialization at $a = b$ reaches a nearby local minimum. By contrast, the fully constrained GD trajectory successfully recovers the ground truth.

Fig. 3b shows distortions of the lattice parameters a and γ of $\text{Na}_3\text{MnCoNiO}_6$, a monoclinic structure with space group Cm (no. 8). Similarly, the unconstrained GD trajectory, even when initialized at a constrained point with $\gamma = 90^\circ$, converges to a local minimum, while the constrained GD trajectory, which enforces $\gamma = 90^\circ$ throughout optimization, reaches the ground truth.

The initialization points for both cases were chosen for illustration, though points for multiple regions would yield similar behavior. These visualizations illustrate how symmetry-constrained XRD-based optimization can better reach the correct phase, yielding higher match rates than its unconstrained counterpart (Fig. 1). This highlights the value of incorporating symmetry constraints into learning schemes that navigate the structure-to-XRD mapping.

Notably, in Fig. S2, we observe roughness that poses challenges for symmetry-constrained GD along symmetry axes such as $a = b$ and $\alpha = 90^\circ$. While fluctuations along $a = b$ are pronounced, those along $\alpha = 90^\circ$ are comparatively shallow and may be mitigated through techniques such as momentum⁴⁰ or regularization, which were not studied in this work. Although symmetry constraints generally improve refinement performance, the landscape visualized here highlights that XRD-based GD may remain sensitive to initialization and prone to local minima in some cases.

In this work, we use an idealized simulated XRD representation (see SI Section 4.7) to isolate the intrinsic properties of the XRD similarity landscape without additional experimental noise. In practice, counting noise in measured diffraction intensities would introduce fluctuations in the similarity metric, further increasing landscape roughness and degrading gradient quality during optimization.

2.1.2 Structure average minimum distances vs. XRD similarity. To assess how effectively diffraction-based similarity metrics capture underlying geometrical similarity, we compare each XRD similarity measure with the Average Minimum Distance (AMD) metric.⁴¹ AMD provides a continuous, geometry-based measure of similarity between periodic

Table 1 Mutual information (MI) between average minimum distance (AMD) and XRD similarity metrics for different noise types and levels

Noise type	Noise level	Cosine	MSE	Entropy	Cross corr.	Cosine + constr.
Lattice	0.01	1.09	0.25	0.76	0.22	1.40
	0.05	0.12	0.16	0.53	0.08	1.55
	0.10	0.09	0.06	0.21	0.07	1.12
	0.15	0.03	0.02	0.02	0.0	0.90
	0.20	0.07	0.02	0.01	0.03	0.85
Coords	0.001	1.05	0.32	0.68	1.02	—
	0.005	0.83	0.48	0.53	0.82	—
	0.010	0.44	0.67	0.56	1.01	—
	0.050	0.96	0.36	1.41	1.81	—
	0.100	0.77	0.32	1.06	1.00	—

structures by computing the Earth Mover's distance between the atomic pointwise distance distributions of the ground-truth and optimized structures. For each structure pair, we compare its AMD with the corresponding XRD similarity score and quantify their relationship using Mutual Information (MI), Pearson correlation, and Spearman correlation.

Correspondence between higher XRD similarity and lower AMD would indicate that diffraction-space similarity aligns with crystal structure similarity. MI captures overall statistical dependence between the two quantities, while Pearson and Spearman correlations describe linear and monotonic relationships, respectively.

As shown in Table 1, the relationship between geometrical and diffraction-based similarity is not uniform across metrics or noise levels. For lattice distortions, cosine similarity yields the highest MI at lowest noise, whereas entropy similarity performs better at moderate noise levels. At higher noise levels, all metrics exhibit similarly low MI, indicating a loss of geometrical correspondence. For coordinate distortions, cosine similarity yields the highest MI at low noise, cross-correlation performs better at moderate noise levels, and entropy performs better at the highest noise level. In contrast, when symmetry constraints are enforced during optimization, MI increases substantially across all lattice noise levels, underscoring the benefit of restricting the search space to symmetry-consistent configurations.

The correlation results reported in Tables S1 and S2 exhibit similar trends. For lattice distortions, no single XRD similarity metric consistently correlates with AMD in the unconstrained setting, whereas enforcing symmetry constraints improves both Pearson and Spearman correlations.

For coordinate distortions, cosine and entropy similarity metrics show negative correlations with AMD, while the MSE-based metric yields a positive and relatively high Spearman correlation. This behavior can be attributed to the fact that MSE more strongly penalizes differences in peak intensities. Since low-level coordinate noise mainly affects peak heights through modifications to the structure factors rather than shifting peak positions, MSE captures these subtle symmetry perturbations more effectively.

Fig. 4 visualizes these relationships. MSE exhibits the weakest correlation with AMD (low Pearson and Spearman






Fig. 4 Average minimum distance (AMD) vs. XRD similarity metrics. (a) Mean squared error (MSE). (b) Entropy similarity. (c) Cross correlation. (d) Cosine similarity. (e) Cosine similarity with lattice constraints applied during optimization. All panels compare structures obtained from XRD-based optimization following lattice distortions of 0.1. Entropy similarity yields higher mutual information (MI) than cosine similarity or MSE at this noise level, though this trend does not generalize across all settings (Table 1). Applying lattice constraints improves MI as well as linear and Spearman correlations.

coefficients), while cosine and entropy similarities show comparable correlations. Entropy similarity achieves slightly higher MI (0.20 vs. 0.09 for cosine), indicating a modestly stronger dependence between diffraction and geometric similarity under these conditions. Nevertheless, the data remain broadly scattered, revealing that all tested similarity metrics struggle to consistently distinguish geometrically distinct configurations. When symmetry constraints are enforced

(Fig. 4d), the correspondence between AMD and cosine similarity improves substantially, both in MI (1.12) and correlation coefficients.

2.2 Metrics used for measuring XRD similarity

Fig. 5 illustrates structure optimization results, with and without lattice constraints. After applying noise, the structure

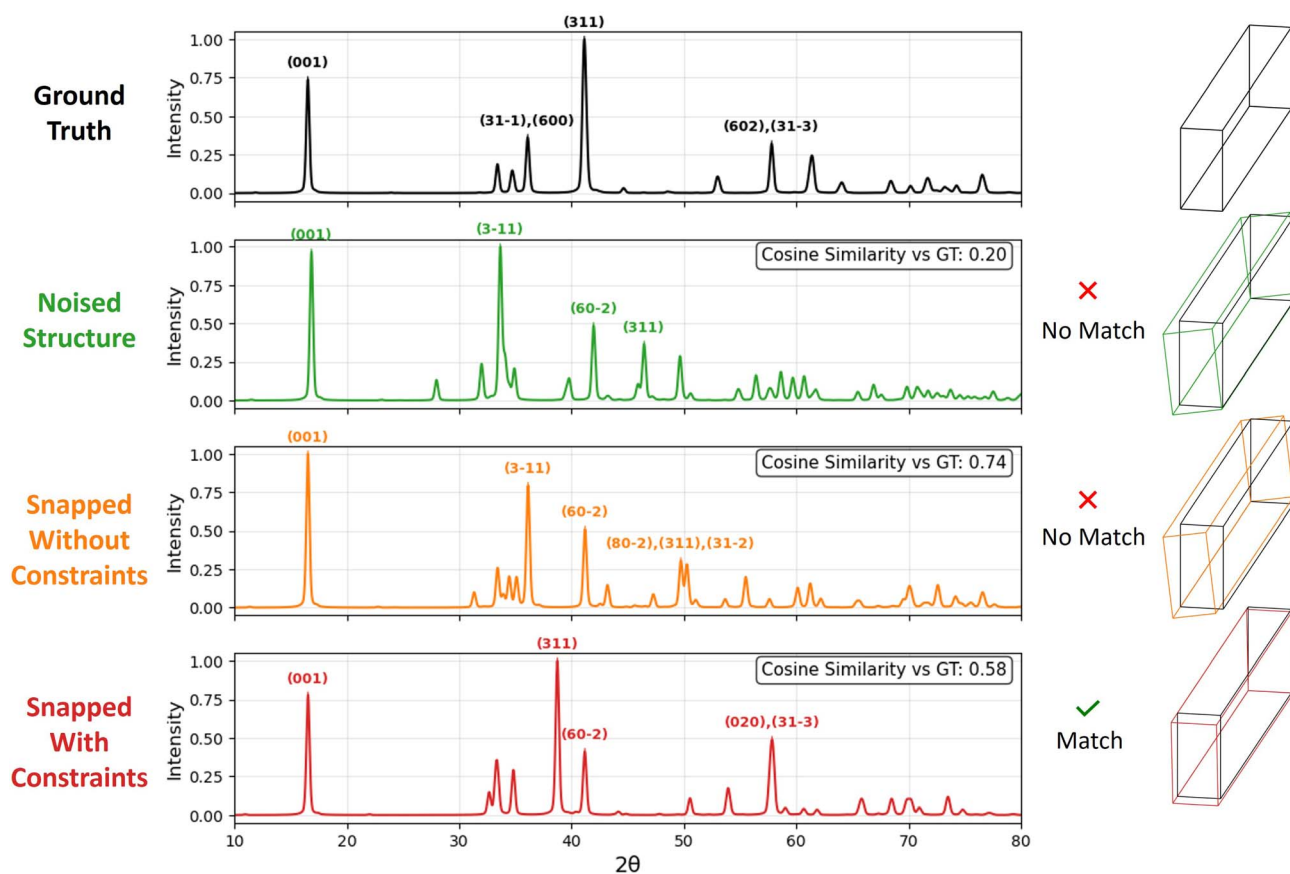


Fig. 5 Lattice and XRD patterns of $\text{Na}_3\text{MnCoNiO}_6$. Each row shows the unit cell relative to the ground truth and corresponding XRD pattern. From top to bottom: ground truth; distorted lattice structure with 0.1 noise level; result of XRD-based GD optimization without constraints; and result of XRD-based GD optimization with symmetry-based constraints. For each, the cosine similarity to the ground truth pattern and the structure match status according to StructureMatcher are reported.



no longer matches the ground truth according to the StructureMatcher metric, with the diffraction pattern exhibiting peak shifts and new reflections. The unconstrained GD optimizer converges to a local minimum where a few peaks align, resulting in a significant increase in cosine similarity (from 0.2 to 0.74). When crystal-family constraints are imposed, the optimizer recovers a structure that matches the ground truth within StructureMatcher tolerance, and the cosine similarity between the patterns increases (from 0.2 to 0.57). However, even in the constrained solution, aligned peaks do not necessarily correspond to identical crystallographic planes, labeled by the Miller indices hkl . *E.g.*, the ground truth 311 plane (top panel) appears slightly shifted in the constrained solution (bottom panel) but aligns with 602, thereby contributing to the similarity score. This discrepancy in interatomic distances seems to remain within the tolerance.

In Fig. S4, the unconstrained GD optimizer converges to a distinct structure whose peaks overlap with those of the ground truth, achieving a deceptively high cosine similarity (0.71). In contrast, the symmetry-constrained optimization successfully recovers a matched structure, but the resulting diffraction pattern shows slightly shifted peaks, leading to a much lower cosine similarity (0.05). The convergence to this

shallow minimum is likely driven by roughness along the symmetry axis, as elaborated in Subsection 2.1.

This counterintuitive outcome highlights a key limitation of using XRD pattern similarity metrics such as cosine similarity, MSE, and entropy similarity as the sole reconstruction objective: similar structures can yield dissimilar patterns, and conversely, distinct structures may appear similar, under such metrics. Because these metrics are insensitive to whether aligned peaks arise from the same atomic geometry, distinct structures can yield high similarity scores, while small geometric deviations can produce low ones. This results in a highly non-convex optimization landscape prone to spurious minima.

This limitation suggests that purely signal-based metrics are unlikely to yield a well-behaved (*e.g.*, convex) optimization landscape. In the absence of smoother metrics, enforcing symmetry-based constraints (see Section 4.4) on the optimizer remains an effective way to navigate otherwise rugged landscapes.

2.3 Comparison to energy relaxation

Structural relaxation through potential energy minimization is often used to refine candidate structures.^{42–46} Fig. 6a shows that the universal ML interatomic potential (MLIP) CHGNet⁴²

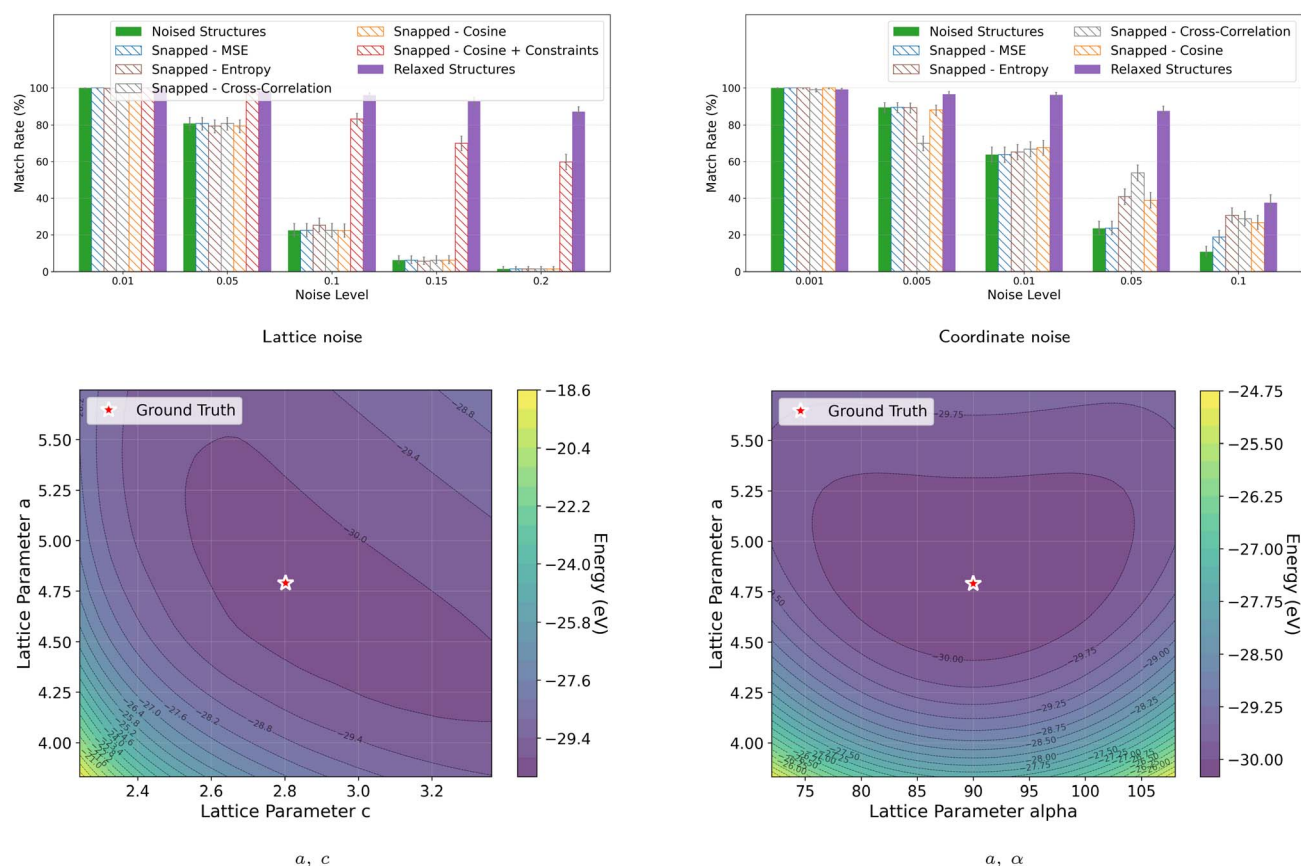


Fig. 6 Comparing XRD-based optimization with energy relaxation. (Top) Match rates from StructureMatcher with ($l_{tol} = 0.1$, $s_{tol} = 0.2$, $angle_{tol} = 5^\circ$) under random lattice and coordinate distortions. Snapped bars are the same as presented in Fig. 1. Energy-based optimization consistently recovers the correct phase, except for high levels of coordinate noise, whereas XRD-based optimization struggles. (Bottom) 2D landscape of CHGNet⁴² predicted energy as a function of lattice parameters of U_2Ti . Along lattice vectors, the energy landscape is smooth and locally convex.



accurately recovers structures matching the ground truth from the distorted state alone, except for a high level (0.1) of coordinate noise (Fig. 6b). In that case, it is plausible that these larger distortions displace the system into the basin of attraction of a different local minimum on the potential energy surface.

These results highlight an important contrast between energy relaxation and XRD-based optimization. Potential-energy landscapes are typically smooth because they arise from local interatomic interactions that vary continuously with atomic positions. Consequently, small perturbations in atomic coordinates generally lead to correspondingly small and continuous changes in the energy.

Cross-sections of the energy optimization landscape are shown in Fig. 6b and S5, which exhibit a smooth and largely convex behavior around the ground-truth structure. This contrast with the highly rugged landscape observed for XRD similarity helps explain why gradient-based relaxation methods are effective for energy minimization but often struggle when applied directly to diffraction-based objectives.

Energy relaxation and XRD-based optimization provide complementary signals: relaxation drives structures toward physically stable configurations, while XRD-based optimization attempts to match an observed pattern. Possible directions for future work can therefore include multi-objective optimization, which can combine the smoothness of the energy landscape with the structure-specific information captured by diffraction.

3 Conclusion

Our analysis suggests that the powder XRD-to-structure mapping presents significant challenges for deterministic, gradient-based optimization. We find that signal-based similarity metrics allow optimizers to converge toward local minima where accidental peak overlaps produce deceptively high similarity scores. We further observe that incorporating symmetry constraints can substantially improve optimization behavior; however, this relies on prior knowledge of the crystal family. This highlights an important aspect of the inverse XRD problem: while diffraction similarity is highly informative when combined with structural priors—as in established refinement approaches—it may not by itself provide a sufficiently well-behaved objective function for reliable structure recovery from arbitrary initializations.

While XRD provides a direct experimental link for generative modeling and the identification of novel phases—a topic of surging interest—current models often produce “near-miss” structures that require further refinement. Our results demonstrate why *post hoc* optimization based solely on XRD similarity is often insufficient for this “last mile” of refinement. While generative models may introduce more complex biases than the physically motivated distortions studied here, our results characterize a loss landscape that is more rugged than the corresponding potential energy surface.

These observations suggest that progress in inverse XRD may be better achieved through generative architectures that condition on XRD patterns while embedding symmetry-aware

inductive biases. In this framework, generative models could provide a robust initialization within the correct structural basin, with final precision guided by the smoother gradients of energy relaxation.

4 Method

10 structures (shown in Fig. S1) were selected from the MP20 dataset,⁴⁷ a collection of small, inorganic, thermodynamically (meta)stable structures derived from the Materials Project in July 2021.⁴⁸ To obtain a representative yet tractable set of systems, we sampled structures from the ten most common space groups in the MP20 dataset, which cover symmetry classes frequently observed in inorganic crystals: $P6/mmm$, $Pn\bar{3}m$, $I\bar{4}$, Cm , $I4/mmm$, $Fm\bar{3}m$, $C2/m$, $P6_3/mmc$, $Pm\bar{3}m$, Pm .

The selected structures also span a range of unit-cell sizes and numbers of atoms while remaining small enough to enable controlled optimization experiments.

For each structure, we generated 50 distorted versions of the conventional cell using two noise models.

4.1 Lattice noise

Random lattice distortions were applied *via* strain tensors.¹¹ These distortions alter the cell while keeping fractional atomic coordinates fixed. Each distorted structure was generated by randomly sampling the entries of a strain tensor and applying it to the ground-truth lattice matrix. Let $\mathbf{L} \in \mathbb{R}^{3 \times 3}$ be the ground-truth lattice matrix (columns are the lattice vectors). The lattice is distorted by a random deformation matrix $\mathbf{S} \in \mathbb{R}^{3 \times 3}$,

$$\tilde{\mathbf{L}} = \mathbf{S}\mathbf{L}, \quad \mathbf{S} = \begin{pmatrix} s_{11} & s_{12} & s_{13} \\ s_{21} & s_{22} & s_{23} \\ s_{31} & s_{32} & s_{33} \end{pmatrix},$$

with no change to the atomic fractional coordinates. For a noise level $\sigma_l > 0$, the entries of \mathbf{S} are sampled as

$$s_{ii} \sim \text{Unif}(1 - \sigma_l, 1 + \sigma_l), \quad i \in \{1, 2, 3\}, \\ s_{ij} \sim \text{Unif}(-\sigma_l, \sigma_l), \quad i \neq j.$$

Thus, diagonal entries produce uniaxial expansion/compression, whereas off-diagonal entries induce shear. By construction, these distortions do not preserve crystal symmetry and can introduce diverse deformation modes. If \mathbf{f} denotes a fractional coordinate, then the Cartesian position changes from $\mathbf{x} = \mathbf{L}\mathbf{f}$ to $\tilde{\mathbf{x}} = \tilde{\mathbf{L}}\mathbf{f}$ with \mathbf{f} unchanged.

4.2 Coordinate noise

Independent, uncorrelated positional displacements were applied to each atom by adding Gaussian-distributed noise to its fractional coordinates. Let $x_i^{(n)} \in [0,1]^3$ be the fractional coordinates of atom n . For a noise scale $\sigma_c > 0$, we draw i.i.d. perturbations

$$\varepsilon_n \sim \mathcal{N}(0, \sigma_c^2 \mathbf{I}_3), \quad n = 1, \dots, N,$$

and set the noisy fractional coordinates to



$$x_f^{(n)} = w(x_i^{(n)} + \varepsilon_n), \quad w(u) = u - \lfloor u \rfloor \in [0,1)^3,$$

where $\lfloor u \rfloor$ applies the floor function component wise. Equivalently, $x_f^{(n)} \equiv x_i^{(n)} + \varepsilon_n \pmod{1}$ (elementwise), *i.e.*, on the 3-torus $\mathbb{T}^3 = \mathbb{R}^3 / \mathbb{Z}^3$.

4.3 Optimization

After noise is applied, gradient-based optimization is performed using StructSnap, a differentiable XRD simulator, with the goal of recovering the ground-truth structure.²⁰ The diffraction patterns are computed from the structure factor contributions of each atomic site, following Bragg's law and the kinematic scattering model. The resulting pattern is a 2D tensor of 2θ angles and intensities.

Either the lattice parameters or atomic coordinates are refined, in accordance with the noise applied. The structure is passed through a differentiable diffraction pipeline to produce a simulated pattern, which is then compared to the ground-truth pattern using a chosen loss function. Given distorted-structure XRD $\hat{\mathbf{x}}$ and ground-truth XRD \mathbf{x} (see SI Section 4.7 for details on XRD calculation), one of the following objective functions is minimized. These include point-wise losses, namely negative cosine similarity, mean squared error (MSE), and negative entropy similarity,³² as well as cross-correlation, which enables shift-based comparison:

$$\mathcal{L}_{\text{cos}} = \frac{1}{N} \sum_{i=1}^N \left(- \frac{\mathbf{x}^{(i)} \cdot \hat{\mathbf{x}}^{(i)}}{\|\mathbf{x}^{(i)}\|_2 \|\hat{\mathbf{x}}^{(i)}\|_2} \right),$$

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N} \sum_{i=1}^N \|\mathbf{x}^{(i)} - \hat{\mathbf{x}}^{(i)}\|_2^2,$$

$$\mathcal{L}_{\text{Entropy}} = -\frac{1}{N} \sum_{i=1}^N \left(1 - \frac{2S_{\hat{x}_i x_i} - S_{x_i} - S_{\hat{x}_i}}{\log 4} \right),$$

$$\mathcal{L}_{\text{XCorr}} = -\frac{1}{N} \sum_{i=1}^N \max_{\tau} \langle \tilde{\mathbf{x}}^{(i)}, \text{shift}_{\tau}(\tilde{\mathbf{x}}^{(i)}) \rangle$$

where S denotes the Shannon entropy and $\tilde{\mathbf{x}} = \frac{\mathbf{x}}{\|\mathbf{x}\|_2}$. Gradients of the loss are backpropagated to update the parameters using PyTorch's autograd.

4.4 Symmetry constraints

The effect of symmetry-based constraints in the lattice-noise case is examined by enforcing the ground-truth crystal family during optimization. At inference time, this ground truth information will not be available. However, previous works have shown strong performance in predicting the crystal family,^{49–52} as well as lattice parameters^{53–55} and space groups^{52,56,57} from XRD patterns. All structures are represented using conventional unit cells, in which crystal-family constraints on the lattice parameters are naturally expressed.

Using projected optimization,⁵⁸ each gradient step is followed by projection onto the constrained values.

Let $\theta = (a, b, c, \alpha, \beta, \gamma)$ be the lattice parameters. Simple constrained gradient descent on a proposed crystal $\hat{\mathcal{X}}$ using

user-selected loss $\mathcal{L} \in \{\mathcal{L}_{\text{cos}}, \mathcal{L}_{\text{MSE}}, \mathcal{L}_{\text{entropy}}\}$ and symmetry projection operator \mathcal{P} would be:

$$\theta^0 = \mathcal{P}(\theta^{\text{init}}), \quad \theta^{k+1} = \mathcal{P}(\theta^k - \nabla_{\theta^k} \mathcal{L}(\text{Xrd}(\hat{\mathcal{X}}(\theta^k)), \text{Xrd}_0)),$$

with the iterations repeating until some convergence criterion is satisfied. Note that θ^{init} are initial lattice parameters from our prediction, model, or, in this case, distorted ground-truth; $\hat{\mathcal{X}}(\theta^k)$ is the representation of our proposed crystal structure, which depends on current lattice parameters θ^k ; Xrd is a map from crystal to computed powder X-ray diffraction pattern; Xrd_0 is the XRD pattern of the reference we aim to recover. The projection operator \mathcal{P} is defined by relevant crystal family, *e.g.*,

$$\mathcal{P}_{\text{cubic}}(a, b, c, \alpha, \beta, \gamma) = (\bar{a}, \bar{a}, \bar{a}, 90^\circ, 90^\circ, 90^\circ), \quad \bar{a} = \frac{a + b + c}{3}.$$

Thus, a , b , and c are first updated independently according to their gradients, then set to the mean value \bar{a} , while angles are fixed to 90° . Similar projectors are defined for the remaining crystal families (see SI Section 4.8 for projector definitions).

4.5 Metrics

Recovery performance is assessed primarily using Match Rate, the fraction of optimized structures that are identified as geometrically equivalent to the ground-truth by StructureMatcher,⁵⁹ considering lattice, atomic positions, and symmetry. The tolerances used are 0.1 for lattice, 0.2 for atomic site positions, and 5° for angles. Additionally, recovery is assessed using the Average Minimum Distance (AMD), which serves as a complementary metric for quantifying geometrical similarity between periodic crystals.⁴¹ While match rate provides a binary classification of whether two structures are equivalent within a given tolerance, AMD offers a continuous measure that compares the distributions of pairwise atomic distances.

Author contributions

N. S. led the implementation, conducted the experiments, performed the analyses, and wrote the manuscript. A. S., M. L., B. K. M., and R. G. B. contributed to experimental conception and provided feedback. B. K. M. and R. G. B. supervised the work. R. G. B. conceived the original idea for the study.

Conflicts of interest

There are no conflicts to declare.

Data availability

An open-source implementation of this work, including the raw data and processing scripts, is available at <https://github.com/learningmatter-mit/XRD-GD> and archived on Zenodo at <https://doi.org/10.5281/zenodo.19336845>.

Supplementary information is available. See DOI: <https://doi.org/10.1039/d6dd00017g>



Acknowledgements

N. S.'s research is sponsored by the Eli and Dorothy Berman Fellowship and the National Science Foundation (NSF) under Award Number 2209892 'Garden: A FAIR Framework for Publishing and Applying AI Models for Translational Research in Science, Engineering, Education, and Industry'.

References

- W. L. Bragg, *Proc. R. Soc. London, Ser. A*, 1914, **89**, 468–489.
- W. I. David and K. Shankland, *Acta Crystallogr., Sect. A: Found. Crystallogr.*, 2008, **64**, 52–64.
- H. A. Hauptman, *Rep. Prog. Phys.*, 1991, **54**, 1427.
- C. Hammond, *The basics of crystallography and diffraction*, Oxford university press, 2015, vol. 21.
- H. M. Rietveld, *J. Appl. Crystallogr.*, 1969, **2**, 65–71.
- S. Gates-Rector and T. Blanton, *Powder Diffr.*, 2019, **34**, 352–360.
- F. Allen and V. Hoy, in *International Tables for Crystallography Volume F: Crystallography of biological macromolecules*, Springer, 2006, pp. 663–668.
- C. Biwer, Z. Feng, D. Finstad, M. McDonnell, M. Knezevic, M. McKerns, D. Savage and S. Vogel, *Sci. Rep.*, 2025, **15**, 8358.
- C. F. Holder and R. E. Schaak, *Tutorial on powder X-ray diffraction for characterizing nanoscale materials*, 2019.
- N. Chandra, K. R. Acharya and P. Moody, *Biol. Crystallogr.*, 1999, **55**, 1750–1758.
- N. J. Szymanski, S. Fu, E. Persson and G. Ceder, *npj Comput. Mater.*, 2024, **10**, 45.
- S. Gu, B. Fu, G. Dodbiba, T. Fujita and B. Fang, *RSC Adv.*, 2017, **7**, 52017–52023.
- H. Deng, K. Wang, Y. Duan, W. Zhang and J. Hu, *Coatings*, 2022, **12**, 429.
- Q. Wang, S. Cai, S. Yang, Y. Yu, Y. Wan, J. Peng, J. Wang and X. Wang, *J. Mater. Sci.: Mater. Electron.*, 2024, **35**, 576.
- A. El-Taher, S. Mansour and I. Lotfy, *J. Mater. Sci.: Mater. Electron.*, 2023, **34**, 599.
- A. Ayeshamariam, M. Kashif, M. Bououdina, U. Hashim, M. Jayachandran and M. Ali, *Ceram. Int.*, 2014, **40**, 1321–1328.
- Z. Ma, D. Li, Y. Han, X. Shi, H. Zhang, D. Song and L. Zhang, *J. Electrochem. Soc.*, 2018, **166**, A5065.
- D. Goranova, E. Lefterova and R. Rashkov, *Int. J. Hydrogen Energy*, 2017, **42**, 28777–28785.
- G. Guo, T. L. Saidi, M. W. Terban, M. Valsecchi, S. J. Billinge and H. Lipson, *Nat. Mater.*, 2025, **1–9**, 1726–1734.
- E. A. Riesel, T. Mackey, H. Nilforoshan, M. Xu, C. K. Badding, A. B. Altman, J. Leskovec and D. E. Freedman, *J. Am. Chem. Soc.*, 2024, **146**, 30340–30348.
- F. L. Johansen, U. Friis-Jensen, E. B. Dam, K. M. Ø. Jensen, R. Mercado and R. Selvan, *arXiv*, 2025, preprint, arXiv:2502.02189, DOI: [10.48550/arXiv.2502.02189](https://doi.org/10.48550/arXiv.2502.02189).
- Q. Li, R. Jiao, L. Wu, T. Zhu, W. Huang, S. Jin, Y. Liu, H. Weng and X. Chen, *Nat. Commun.*, 2025, **16**, 7428.
- Q. Lai, F. Xu, L. Yao, Z. Gao, S. Liu, H. Wang, S. Lu, D. He, L. Wang, L. Zhang, *et al.*, *Adv. Sci.*, 2025, **12**, 2410722.
- G. Guo, J. Goldfeder, L. Lan, A. Ray, A. H. Yang, B. Chen, S. J. Billinge and H. Lipson, *npj Comput. Mater.*, 2024, **10**, 209.
- K. Choudhary, *J. Phys. Chem. Lett.*, 2025, **16**, 2110–2119.
- A. S. Parackal, R. E. Goodall, F. A. Faber and R. Armiento, *Phys. Rev. Mater.*, 2024, **8**, 103801.
- J. Lee, J. Oba, N. Ohba and S. Kajita, *npj Comput. Mater.*, 2023, **9**, 135.
- Y. Gevorkov, O. Yefanov, A. Barty, T. A. White, V. Mariani, W. Brehm, A. Tolstikova, R.-R. Grigat and H. N. Chapman, *Acta Crystallogr., Sect. A: Found. Crystallogr.*, 2019, **75**, 694–704.
- A. Otero-de-la Roza, *Appl. Crystallogr.*, 2024, **57**, 1401–1414.
- S. Racioppi, A. Otero-de-la Roza, S. Hajinazar and E. Zurek, *Digital Discovery*, 2025, **4**, 73–83.
- E. Hernández-Rivera, S. P. Coleman and M. A. Tschopp, *ACS Comb. Sci.*, 2017, **19**, 25–36.
- Y. Li, T. Kind, J. Folz, A. Vaniya, S. S. Mehta and O. Fiehn, *Nat. Methods*, 2021, **18**, 1524–1531.
- O. Cannelli, J. Wiktor, N. Colonna, L. Leroy, M. Puppini, C. Bacellar, I. Sadykov, F. Krieg, G. Smolentsev, M. V. Kovalenko, *et al.*, *J. Phys. Chem. Lett.*, 2022, **13**, 3382–3391.
- F. Brivio, J. M. Frost, J. M. Skelton, A. J. Jackson, O. J. Weber, M. T. Weller, A. R. Goni, A. M. Leguy, P. R. Barnes and A. Walsh, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 2015, **92**, 144308.
- L. Delgado-Aparicio, M. Bitter, Y. Podpaly, J. Rice, W. Burke, M. S. Del Rio, P. Beiersdorfer, R. Bell, R. Feder, C. Gao, *et al.*, *Plasma Phys. Controlled Fusion*, 2013, **55**, 125011.
- R. Jiao, W. Huang, Y. Liu, D. Zhao and Y. Liu, *The Twelfth International Conference on Learning Representations*, 2024.
- D. Levy, S. S. Panigrahi, S.-O. Kaba, Q. Zhu, K. L. K. Lee, M. Galkin, S. Miret and S. Ravanbakhsh, *arXiv*, 2025, preprint, arXiv:2502.03638, DOI: [10.48550/arXiv.2502.03638](https://doi.org/10.48550/arXiv.2502.03638).
- N. Kazeev, W. Nong, I. Romanov, R. Zhu, A. Ustyuzhanin, S. Yamazaki and K. Hippalgaonkar, *arXiv*, 2025, preprint, arXiv:2503.02407, DOI: [10.48550/arXiv.2503.02407](https://doi.org/10.48550/arXiv.2503.02407).
- L. D. Brown, T. T. Cai and A. DasGupta, *Stat. Sci.*, 2001, **16**, 101–133.
- N. Qian, *Neural Networks*, 1999, **12**, 145–151.
- D. Widdowson and V. Kurlin, *Adv. Neural Inf. Process. Syst.*, 2022, **35**, 24625–24638.
- B. Deng, P. Zhong, K. Jun, J. Riebesell, K. Han, C. J. Bartel and G. Ceder, *Nat. Mach. Intell.*, 2023, **5**, 1031–1041.
- I. Batatia, P. Benner, Y. Chiang, A. M. Elena, D. P. Kovács, J. Riebesell, X. R. Advincula, M. Asta, M. Avaylon and W. J. Baldwin, *J. Chem. Phys.*, 2025, **163**, 184110.
- C. Chen and S. P. Ong, *Nat. Comput. Sci.*, 2022, **2**, 718–728.
- A. Musaelian, S. Batzner, A. Johansson, L. Sun, C. J. Owen, M. Kornbluth and B. Kozinsky, *Nat. Commun.*, 2023, **14**, 579.
- B. M. Wood, M. Dzamba, X. Fu, M. Gao, M. Shuaibi, L. Barroso-Luque, K. Abdelmaqsoud, V. Gharakhanyan, J. R. Kitchin, D. S. Levine, *et al.*, *arXiv*, 2025, preprint, arXiv:2506.23971, DOI: [10.48550/arXiv.2506.23971](https://doi.org/10.48550/arXiv.2506.23971).
- T. Xie, X. Fu, O.-E. Ganea, R. Barzilay and T. S. Jaakkola, *International Conference on Learning Representations*, 2021.



- 48 A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, *et al.*, *APL Mater.*, 2013, **1**, 1.
- 49 Y. Suzuki, H. Hino, T. Hawaii, K. Saito, M. Kotsugi and K. Ono, *Sci. Rep.*, 2020, **10**, 21790.
- 50 B. D. Lee, J.-W. Lee, W. B. Park, J. Park, M.-Y. Cho, S. Pal Singh, M. Pyo and K.-S. Sohn, *Adv. Intell. Syst.*, 2022, **4**, 2200042.
- 51 S. Zhang, B. Cao, T. Su, Y. Wu, Z. Feng, J. Xiong and T.-Y. Zhang, *IUCrJ*, 2024, **11**, 634–642.
- 52 C. Bin, Y. Liu, Z. Zheng, R. Tan, J. Li and T.-y. Zhang, *The Thirteenth International Conference on Learning Representations*, 2025.
- 53 S. Habershon, E. Y. Cheung, K. D. Harris and R. L. Johnston, *J. Phys. Chem. A*, 2004, **108**, 711–716.
- 54 S. R. Chitturi, D. Ratner, R. C. Walroth, V. Thampy, E. J. Reed, M. Dunne, C. J. Tassone and K. H. Stone, *Appl. Crystallogr.*, 2021, **54**, 1799–1810.
- 55 H. Dong, K. T. Butler, D. Matras, S. W. Price, Y. Odarchenko, R. Khatry, A. Thompson, V. Middelkoop, S. D. Jacques, A. M. Beale, *et al.*, *npj Comput. Mater.*, 2021, **7**, 74.
- 56 H. Schopmans, P. Reiser and P. Friederich, *Digital Discovery*, 2023, **2**, 1414–1424.
- 57 F. Oviedo, Z. Ren, S. Sun, C. Settens, Z. Liu, N. T. P. Hartono, S. Ramasamy, B. L. DeCost, S. I. Tian, G. Romano, *et al.*, *npj Comput. Mater.*, 2019, **5**, 60.
- 58 S. Wright, J. Nocedal, *et al.*, *Springer Sci.*, 1999, **35**, 7.
- 59 S. P. Ong, W. D. Richards, A. Jain, G. Hautier, M. Kocher, S. Cholia, D. Gunter, V. L. Chevrier, K. A. Persson and G. Ceder, *Comput. Mater. Sci.*, 2013, **68**, 314–319.

