

Dara: Automated multiple-hypothesis phase identification and refinement from powder X-ray diffraction

Yuxing Fei,^{†,‡,§} Matthew J. McDermott,^{‡,§} Christopher L. Rom,[¶] Shilong Wang,^{†,‡} and Gerbrand Ceder*,^{†,‡}

[†]*Department of Materials Science & Engineering, University of California, Berkeley, CA 94720, USA*

[‡]*Materials Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA*

[¶]*Materials Science Center, National Renewable Energy Laboratory, Golden, CO, 80401, USA*

[§]*Authors contribute equally*

E-mail: gceder@berkeley.edu

Abstract

Powder X-ray diffraction (XRD) is a foundational technique for characterizing crystalline materials. However, the reliable interpretation of XRD patterns, particularly in multiphase systems, remains a manual and expertise-demanding task. As a characterization method that only provides structural information, multiple reference phases can often be fit to a single pattern, leading to potential misinterpretation when alternative solutions are overlooked. To ease humans' efforts and address the challenge, we introduce Dara (Data-driven Automated Rietveld Analysis), a framework designed to automate the robust identification and refinement of multiple phases from powder XRD data. Dara performs an exhaustive tree search over all plausible phase combinations within a given chemical space and validates each hypothesis using a robust Rietveld refinement routine (BGMN). Key features include structural database filtering, automatic clustering of isostructural phases during tree expansion, peak-matching-based scoring to identify promising phases for refinement. When ambiguity exists, Dara generates multiple hypothesis which can then be decided between by human experts or with further characterization tools. By enhancing the reliability and accuracy of phase identification, Dara enables scalable analysis of realistic complex XRD patterns and provides a foundation for integration into multimodal characterization workflows, moving toward fully self-driving materials discovery.

Introduction

Accurate characterization of material structures is essential for understanding synthesis-structure-property relationships in materials. For bulk inorganic materials, powder X-ray diffraction (XRD) has long been a pivotal and widely used technique for determining majority-phase crystal structures.^{1,2} Moreover, powder XRD plays a critical role in the discovery of inorganic materials, serving as a key tool for confirming the synthesis of predicted target structures, such as those derived from *ab initio* methods.^{3,4} With appropriate analysis, XRD patterns can provide valuable insights into material properties, including phase fractions, lattice parameters, strains, site occupancies, and more.⁵

XRD analysis typically begins with the identification of all phases present in the pattern. This process involves comparing the experimental pattern with the calculated patterns of structures available in crystal structure databases like the Materials Project (MP),⁶ Inorganic Crystal Structure Database (ICSD),⁷ Powder Diffraction File (PDF),⁸ and Crystallography Open Database (COD).⁹ The task becomes practically challenging when a sample contains multiple phases that cannot be perfectly matched to reference structures. This is common when characterizing the synthesis products of exploratory inorganic reactions or natural minerals, which may exhibit compositional variance or preferred orientation effects. In these cases, accurate interpretation requires meticulous analysis and the expertise of researchers who are intimately familiar with the material system, enabling them to discern its subtle nuances.

In recent years, the development of automated and autonomous laboratories for the discovery of inorganic materials has underscored the need to accelerate the characterization process.^{10–13} As the throughput of synthesis and characterization continues to increase, human analysis of patterns has become impractical, further emphasizing the importance of automation. The integration of reliable, robust, and accurate powder XRD phase identification algorithms within autonomous laboratories will be crucial to establishing high-quality experimental databases for inorganic materials. In self-driving labs, high-quality character-

ization of samples at the early stage of synthesis optimization is of particular importance in the AI-driven decision-making algorithms, which may be more challenged by erroneous interpretations than human experts are.

Algorithms for XRD phase analysis have captured the interest of researchers for nearly a century. In 1938, Hanawalt¹⁴ introduced a qualitative method for phase identification based on major diffraction peaks of known phases. Using this approach, researchers can refer to the “Hanawalt Index”, a tabulated collection of peak data, to identify potential phases. If three major peaks align with those of a specific phase, it strongly suggests the presence of that phase in the sample. This method is also known as the “search-match” method, as it always involves searching the reference database and then matching the phase to the diffraction pattern. Following Hanawalt’s manual search, several computer programs are available for automated peak indexing,^{15–17} each with carefully fine-tuned strategies to identify phases. With more peaks taken into consideration and user-friendly interfaces, these programs enable researchers to analyze XRD patterns with rigor and ease. This makes them the dominant method for phase analysis due to their straightforward nature and low computational resource requirements.

With the growing trend of applying deep learning to XRD analysis, numerous studies have leveraged neural networks (NNs) to automate the interpretation of diffraction patterns. Oviedo et al.¹⁸ used a convolutional neural network (CNN) with simulated and experimental XRD data to classify crystal dimensionality and space group, achieving 93% and 89% accuracy, respectively, by augmenting limited data with physics-informed simulations. Lee et al.¹⁹ tackled multi-phase identification by training a deep CNN on over 1.7 million synthetic mixed XRD patterns (combinations of 170 compounds in the Sr-Li-Al-O quaternary system), enabling near-100% phase identification in complex mixtures and even quantifying phase fractions with 86% accuracy. Maffettone et al.²⁰ employed an ensemble of CNNs (a “crystallography companion agent”) that outputs probabilistic phase predictions, avoiding combinatorial explosion in training while providing confidence metrics for each identified

phase. Szymanski et al.²¹ combined physics-informed peak perturbation augmentation with an ensemble CNN and a branching algorithm to iteratively identify phases in mixtures. This probabilistic approach achieved higher accuracy than traditional profile-matching and earlier deep-learning methods on challenging multi-phase samples. More recently, researchers have created a new NN architecture for the XRD phase identification task. For example, Zhang et al.²² introduced a self-attention CNN (CPICANN) trained on around 700k simulated patterns (from 23k structures), which attained 98.5% accuracy on single-phase identification with element information provided and 80% on experimental scans, significantly outperforming conventional XRD software. Beyond deep NNs, other machine-learning approaches have also been explored. For example, Suzuki et al.²³ used an interpretable tree-ensemble model to classify crystal systems and space groups with 90% accuracy, revealing human-understandable diffraction features.

Alongside computerized Hanawalt methods and their NN-based variations, full-profile search-match methods have recently gained attention. These methods use pattern-fitting programs to match calculated patterns to experimental data. For example, Lutterotti et al.²⁴ used results from Rietveld refinement to calculate a figure of merit (FoM) that measures phase fitness, considering factors such as the refinement R-value, density differences, crystallite size, and microstrain. Chang et al.²⁵ proposed a pseudo-refinement method called CrystalShift, which uses a best-first tree search to refine phase combinations and applies Bayesian model comparison to estimate phase probabilities without requiring additional phase space information or training. These methods can use more detailed peak models to handle patterns with poor crystallinity or highly oriented grains. They also improve interpretability by providing additional information through the refinement process, helping to better understand the fitness of the output phase combination.

Despite the numerous approaches proposed for reliable phase identification in XRD, they still occasionally yield inaccurate results. This limitation arises from the nature of XRD as a structure-based technique, which inherently lacks information about the composition of

phases.^{26,27} A diffraction pattern can often be interpreted as different combinations of reference phases due to the presence of isostructural phases and solid solutions in the structure database. Furthermore, when a sample contains a mixture of phases, the minor peaks of some phases may overlap with the dominant peaks of others. This peak overlap can result in several plausible phase combinations that fit the pattern, making it impossible to definitively determine the correct solution based solely on XRD data without external knowledge of the material system.

Given these challenges, a reliable automated XRD analysis tool should be able to present multiple possible phase combinations matching a pattern when ambiguity exists. For example, when characterizing a synthesized sample of a purported solid solution, the tool should be able to present and compare the null hypothesis (a multi-phase combination of end-members) with the desired result (a single-phase solid solution). This is especially critical whenever the endmembers are isostructural with the target solid solution. Moreover, in cases where the null or alternative hypotheses have a similar quality of fit, the tool should be able to provide hints for further characterization to disambiguate different solutions.

Motivated by the need to address the ambiguity issue in automated XRD pattern analysis, we present Dara, a data-driven Rietveld analysis framework. Dara is designed to generate all validated hypotheses that align well with a given XRD pattern, offering a comprehensive and reliable solution for phase identification. Dara employs an exhaustive tree search algorithm complemented by intelligent composition and structure grouping strategies to achieve robust XRD phase identification and ensure human readability. It is particularly suited for analyzing the phases in complex, multi-phase diffraction patterns, such as those obtained from powder products of solid-state reactions. Unlike the conventional use of Rietveld refinement, which primarily extracts phase structure information out of XRD patterns, Dara leverages refinement engines (e.g., BGMN) earlier in the analysis pipeline to identify candidate phases, ensuring both transparency and interpretability. Nonetheless, Dara is not intended for detailed structural refinement, such as determining atomic positions, site oc-

cupancies, or displacement parameters, which require expert knowledge and task-specific context.

The features of Dara include:

- Ability to analyze complex, multi-phase diffraction patterns of solid-state reaction products, primarily focusing on phase identification
- Rietveld-refinement-based search algorithm that can provide good interpretability.
- Null hypothesis generation and testing engine to explore all possible combinations of phases. If multiple combinations fit well, they are all provided and ranked.
- Compatible with reference structures for multiple sources, supporting structures from both experimental and computational databases.
- Designed for future integration in multi-modal characterization efforts: XRD and other elemental analyses like SEM/EDS, XRF, and XPS.

Methods

The XRD analysis workflow of Dara is shown in Figure 1, with details of each step described in the following subsections. In brief, the analysis begins with a set of reference phases, which are a list of crystalline material structures. These may be supplied by the user or generated automatically by Dara by filtering database entries to the sample’s element set and removing redundant entries (Figure 1(a)). Dara then constructs a search tree to iteratively explore the likely reference phase combinations (Figure 1(b)) and identify all phases that can be present in a sample, potentially containing multiple phases, from the reference phase set. Each node represents a phase combination, which is a subset of the reference phase set. A directed edge adds one phase to a node’s phase combination, producing a child node. Because exhaustive traversal of all combinations is combinatorial and intractable, Dara employs a heuristic peak-matching score, similar to the search-match method, to prioritize promising phases

and prune unlikely phases (Figure 1(c)). A Rietveld refinement then evaluates the phase combination in a node to obtain accurate fit metrics (Figure 1(d)). After that, Dara will expand the node by adding one more phase, creating a new node in the tree. The search ends either when a user-defined maximum number of phases in a combination is reached or when adding phases no longer improves the fit. Finally, Dara retains well-fitting phase combinations, discards poorly fitted phase combinations, and groups phase combinations by diffraction similarity and composition to present an interpretable summary of the likely phase combinations in the sample (Figure 1(e)).

Structural databases & preprocessing

Dara’s analysis workflow starts with a set of reference phases, which are known crystalline structures, typically observed experimentally. Users can either supply the reference phases on their own, or use Dara’s workflow to automatically generate and clean up the reference phases for an input XRD pattern. The workflow begins with input structure databases (e.g., COD), which undergo a series of preprocessing steps to filter redundant, problematic, and high-energy phases before being passed to the downstream tree search. Although Dara can work with reference phases without filtering, this pre-processing workflow specifically removes molecular, organic, and duplicate phases, consistent with the typical target application of analyzing inorganic crystalline solids in powder diffraction patterns.

Duplicate phases are first identified via the structure matching algorithm in the pymatgen package.²⁹ From each duplicate set, we select the phase characterized at a temperature closest to 20°C and recorded earliest (i.e., with the oldest entry year) in the database. For the cases where the XRD pattern is not measured at room temperature, such as those acquired during *in situ* heating, we find that the variations in lattice parameters and Debye-Waller factors are typically minor enough for Dara to accurately identify the correct phases.

To filter down the reference phases to the most plausible set, the DFT energies are retrieved from the Materials Project (MP) database.⁶ Because exact matches between ex-




Figure 1 Overview schematic of XRD phase analysis performed with Dara. (a) The pre-processing workflow filters reference phases, which are a set of crystalline material structures, from structure databases such as the Crystallography Open Database (COD).^{9,28} First, all phases within the chemical system of the input XRD pattern are selected. Then, duplicate phases are removed based on the formula and space group. High-energy phases are further filtered out using thermodynamic data from the Materials Project. The resulting phases are used as the reference phases in the downstream search routine. (b) A search tree is constructed with each node representing a phase combination, and directed edges representing the addition of one phase to the previous node's phases. The color of each node represents the weighted profile residual (R_{wp}) values. Darker colors represent lower R_{wp} (indicating a better fit). (c) A peak matching algorithm to quickly filter phases that can fit well to any of the remaining unmatched peaks to prune unlikely phases and save computation time. (d) Identified phases are then passed to a Rietveld refinement engine, such as BGMN. The black crosses are the experimental pattern. The orange line represents the calculated pattern output by Rietveld refinement. (e) Multiple results are extracted from the search tree and presented to the user. The results are ranked by R-values and grouped based on their compositions and structures for easier interpretation. Results with excessively high R-values are removed.

perimental and DFT-optimized structures are not always available, the energy of the lowest-energy MP structure with the same space group and composition as the experimental phase is used as an approximation. Phases with an energy above the hull greater than 100 meV/atom are discarded, while phases without a corresponding MP entry are retained to avoid inadvertently excluding them.

Tree search for phase identification

Once the reference phase set is established, Dara initiates its core analysis by constructing a search tree. The search tree systematically explores all viable combinations of reference phases (Figure 1(b)). Each node in the tree corresponds to a specific phase combination, while each directed edge represents the addition of one new phase. This process, known as node expansion, generates new phase combinations by adding one phase at a time to an existing node, thereby forming a new child node. For each node, Dara will perform Rietveld refinement to evaluate the fitness of phases to the patterns. Because the refinement is computationally expensive, Dara restricts expansion to a shorter list of promising reference phases, selected by scoring all remaining references in the chemical system with the peak-matching algorithm described later. To avoid redundant exploration of the same phase combinations in different orders across the search tree, Dara enforces an ordering constraint during tree expansion. Each newly added phase must have a lower maximum peak intensity than that of any phase already in the node. This ensures that each unique phase combination is visited only once, progressing from the most prominent phases to the less prominent ones.

In XRD phase analysis, it is common to encounter multiple reference phases that share almost identical diffraction patterns, such as solid solutions with slight variations in composition or those with minor differences in atomic orderings. These phases almost always yield a similar fit during refinement. It is advantageous to group these phases and treat them collectively. Dara implements this strategy by using the peak-matching algorithm to quantify the similarity of diffraction patterns (described in *Peak matching* section). Furthermore, to

avoid redundant refinements within a group, Dara only continues node expansion with one representative phase per group. The representative phase is chosen using the figure of merit (FoM),^{24,30} which considers the quality factor ($1 - \rho$, an R-factor variation defined in BGMN to eliminate background effect) and the lattice parameter shift (ΔU). A detailed explanation of the quality factor can be found in the BGMN user guide.³¹ The FoM is calculated as

$$\text{FoM} = \frac{1}{(1 - \rho) + \Delta U} \quad (1)$$

$$\Delta U = 100 \cdot \left(\frac{|a_{\text{refined}} - a_0|}{a_0} + \frac{|b_{\text{refined}} - b_0|}{b_0} + \frac{|c_{\text{refined}} - c_0|}{c_0} \right),$$

where $\{a, b, c\}_{\text{refined}}$ represent the lattice parameters obtained from refinement, and $\{a, b, c\}_0$ represents the lattice parameters of the unrefined reference phase. The algorithm prioritizes phases that require smaller lattice parameter shifts during refinement, thereby avoiding the overfitting of solid solutions. To ensure the newly added phase improves the fit, Dara evaluates the improvement in the profile residual factors with corrected background (R_{pb}),³¹ an R-factor that eliminates the background effect by subtracting the fitted background. If the improvement in R_{pb} falls below a specified threshold, the node will no longer be expanded. Otherwise, the search continues until a maximum number of phases is reached. In Dara, a default R_{pb} improvement threshold of 2% is used to balance overfitting and underfitting of the pattern. The default maximum number of phases is set to 5, as peak overlap typically becomes too severe to distinguish phases beyond this point reliably. At the end of the search routine, all grouped phases, including those with lower FoM, are reported together for user consideration.

To accelerate the tree search process, Dara utilizes the Ray framework³² to run multiple tree node expansions and Rietveld refinements concurrently. A breadth-first search (BFS) strategy is employed, with an internal queue managing node expansion tasks. Worker processes pull tasks from the queue, with each initiating a subtree expansion that includes peak matching and Rietveld refinement. Upon completion, the expanded subtree is merged back

into the main search tree managed by the master process, which then identifies and queues new nodes for further expansion. Thanks to the high scalability of Ray, Dara can run efficiently on multi-core CPUs and even across multiple nodes in high-performance computing (HPC) clusters, significantly reducing analysis latency.

Peak matching

During node expansion, Dara performs a Rietveld refinement on each proposed phase combination. However, as the size of the reference phase set increases, the number of possible combinations, and thus refinements, would grow exponentially. To avoid unnecessary refinements, Dara uses a customized peak-matching algorithm to quickly estimate the fitness of a phase to the XRD pattern before committing to a Rietveld refinement. This algorithm finds a mapping between peaks in an experimental pattern (measured on the actual sample) and those in the calculated XRD pattern of each reference phase. Then, a heuristic fitness score is computed based on the mapping and used to evaluate the fitness of a phase to the experimental pattern.

Before constructing the search tree, Dara extracts all the peaks in the experimental pattern using the TEIL&EFLECH³³ program within the BGMN software suite, which can accommodate patterns with significant peak overlap. To obtain the calculated peaks in each reference phase, Dara also performs a single-phase refinement for every reference phase. The goal of this step is to generate calculated diffraction peaks for each phase, which serve as the basis for comparison with the experimental pattern during peak matching. Peaks in both experimental and calculated patterns are classified into four categories (Figure 1(c)): matched, wrong intensity, missing, and extra. The matched peaks refer to peaks that appear in both the calculated and experimental patterns at nearly the same position and with similar intensities. The wrong intensity peaks refer to those that appear in similar positions but with very different intensities (i.e., by a factor greater than five, which is the default used in Dara). The missing peaks refer to those that appear only in the experimental pattern but

not in the calculated pattern, which can indicate either a poor fit or the existence of other phases. The extra peaks refer to those that appear only in the calculated pattern but not in the experimental pattern. This sometimes occurs when the Rietveld refinement procedure determines that it is a mathematically more optimal fit to match only the major peaks in an experimental pattern while leaving some minor peaks unmatched. Extra peaks in the calculated pattern typically indicate a structural discrepancy, such as symmetry breaking from different atom ordering, between the actual material and the reference phase, making the latter less favorable for selection in phase identification. After classification, a heuristic score for each reference phase is calculated based on the intensity of peaks in different categories, as

$$\text{Score} = \frac{\sum I_{\text{matched}} + \sum I_{\text{wrong intensity}} - 0.1 \sum I_{\text{missing}} - 0.5 \sum I_{\text{extra}}}{\sum I_{\text{exp}}}, \quad (2)$$

where I is the intensity of each peak, with I_{matched} , $I_{\text{wrong intensity}}$, and I_{missing} referring to the intensity of matched, wrong-intensity, and missing peaks in the experimental peak list, respectively. $I_{\text{extra,calculated}}$ refers to the extra peaks in the calculated peak list. The score is normalized by I_{exp} , the total intensity of all experimental peaks. The score function is designed such that the presence of missing and extra peaks penalizes (i.e., decreases) the score. In contrast, a greater number of matched and wrong-intensity peaks increases the score. The coefficients were determined through a heuristic process and can be adjusted.

By calculating scores for all reference phases, Dara can quickly identify phases that potentially have a good fit and warrant further Rietveld refinement, which is a slower but more accurate process. When a node is expanded during the tree search, the missing experimental peaks are extracted by peak matching algorithm, indicating none of the phases in the node can account for these peaks. Afterwards, a new peak matching is performed for each reference phases' calculated peaks against the missing experimental peaks. The score is then calculated from the peak matching result to measure the fitness of a reference phase. In prac-

tice, only a small subset of phases achieves high scores, indicating good fitting. The majority show poor agreement due to the significant difference between the calculated pattern and the experimental XRD pattern. Instead of applying a fixed threshold, Dara dynamically determines a threshold by detecting the transition between high- and low-scoring phases. This is achieved by analyzing the cumulative percentile distribution of scores. The maximum of its second derivative (the inflection point) marks where the scores shift most sharply from good to poor, which can be seen as a boundary between good and poor fitting. Only phases with a score higher than this threshold will be added to the phase combination to create new nodes, where Rietveld refinement is performed to evaluate the fitness more accurately.

In addition to quickly estimating the fitness of a phase, Dara also employs the peak matching algorithm to group phases with similar XRD patterns during node expansion. To this goal, Dara runs the aforementioned peak-matching algorithm between two calculated patterns from two different reference phases and classifies each peak into one of the four categories. Then, a Jaccard index is used to calculate the similarity between the two computed patterns, as

$$\text{Similarity}(\text{Pattern 1}, \text{Pattern 2}) = \frac{\sum I_{\text{matched+wrong intensity}}^{\text{Pattern 1}} + \sum I_{\text{matched+wrong intensity}}^{\text{Pattern 2}}}{\sum I^{\text{Pattern 1}} + \sum I^{\text{Pattern 2}}}. \quad (3)$$

For each node expansion, the pairwise similarity between all newly added phases is calculated, forming a similarity matrix. Then, an agglomerative clustering algorithm³⁴ is applied to the similarity matrix to group phases that exhibit nearly identical calculated patterns. We use the clustering algorithm implemented in scikit-learn,³⁵ with a default similarity threshold of 0.9. Only one representative node in each node group is selected with FoM and continues the node expansion. Others will be considered as alternative structure solution, as described in the *Tree search for phase identification* section.

BGMN Rietveld refinement engine

Rietveld refinement in Dara is performed with the BGMN^{36,37} package, as shown in Figure 1(d). The BGMN binaries (v4.2.23) are used as compiled and supplied by the Profex team.³⁰ Two sets of refinement parameters are used in Dara: (1) search refinement parameters, which are used during the phase search stage and restrict peak broadening and preferred orientation to avoid overfitting, and (2) final refinement parameters, which allow a wider range of adjustments (e.g., larger peak broadening and preferred orientation).

For the refinement parameters in the phase search stage, a maximum of 1% lattice strain is allowed on each phase during refinement. A peak model with Cauchy square broadening ($r_p=4$) is used to describe the peak shape. The width parameter, k_1 , is constrained between 0 and 0.01. The second width parameter, k_2 , is fixed to 0. The crystalline size parameter, B_1 , is constrained between 0 and 0.005. The weight fraction of each phase is calculated directly by normalizing the scale factor (GEWICHT) of all the phases in each pattern. The sample displacement factor, EPS2, is constrained between -0.05 and 0.05.

After the phase search stage, a finer refinement step is conducted on each searched phase combination to obtain a more accurate fit of the XRD peaks and determine phase fractions. In this refinement, up to 1% lattice strain is allowed for each phase. Peak shapes are modeled using a Cauchy-squared broadening function with a profile parameter $r_p = 4$. The peak width parameter k_1 is constrained between 0 and 0.01, while k_2 is fixed at 0. The crystallite size parameter B_1 is constrained between 0 and 0.05. Preferred orientation is accounted for using a fourth-order spherical harmonic function (SPHAR4). Phase weight fractions are calculated by normalizing the scaling factor (GEWICHT) across all identified phases in a given pattern. The sample displacement parameter EPS2 is constrained between -0.05 and 0.05 to correct for sample height effects.

Result representation and compositional grouping

At the end of the tree search, Dara collects all leaf nodes. These phase combinations can span a wide range of fit quality, from poor to good, so further filtering is required. To retain only meaningful results, Dara applies the Jenks natural break detection algorithm,³⁸ which clusters phase combinations with similar quality factor ($1 - \rho$) together by minimizing variance of quality factors within each cluster. Dara then reports the cluster of solutions with the lowest quality factor as the final result of phase combinations that can fit the pattern while discarding those with a clearly poor fit.

For easier human interpretation, the identified phases are also further grouped by composition. This compositional grouping is performed using an agglomerative clustering algorithm.³⁴ For each composition group, the representative composition is chosen as the integer composition closest to the average composition. If no integer composition is available, the selected composition is the one closest to the average composition of the group. By clustering compositionally similar phases, this approach reduces redundancy in the reported results and highlights the most relevant, representative phases.

Additionally, thanks to the peak-matching algorithm, the search results include information describing any unmatched peaks. If a given phase combination cannot account for all the peaks in the experimental pattern or introduces extra peaks, the peak-matching algorithm flags these discrepancies. This provides an additional measure of the quality of phase identification, allowing users to assess the extent to which the proposed phase combination accurately explains the experimental data. The reported differences serve as indicators of potential missing or misidentified phases, guiding further refinement in phase selection and structure analysis.

Results

To evaluate Dara’s performance in analyzing real-world XRD patterns, we create several benchmark scenarios of real powder diffraction patterns. The results of these tests are described in the following sections.

Benchmarking on a dataset of commercial precursor mixtures

For the first test case, we construct a benchmark XRD pattern dataset by mixing commercial precursor materials (oxides and carbonates) in varying ratios. Ten crystalline, single-phase precursors are selected and used to prepare 10 binary (two-component) and 10 ternary (three-component) mixtures. In each mixture, the precursors are randomly selected to constitute between 10 wt% and 90 wt% of the total mass, yielding a wide range of peak intensities to assess Dara’s performance across different weight fractions, as illustrated in Figure 2(a). Detailed preparation protocols are provided in Supplementary Note ???. Each sample is measured using two scan settings: a short scan (2 minutes, low quality) and a longer scan (8 minutes, medium quality) between 10° and 100° , enabling evaluation across different noise levels.




Figure 2 Preparation of the precursor mixture dataset and benchmarking results. (a) Schematic illustration of the procedure for generating the precursor mixture dataset. Ten commercial precursors are randomly selected and mixed at varying ratios between 10 wt% and 90 wt%, resulting in 10 binary and 10 ternary precursor mixtures. XRD patterns are collected using a benchtop diffractometer under two scanning programs (2 minutes and 8 minutes) to produce datasets of different measurement qualities. (b) Comparison of correctly indexed patterns by Jade and Dara. Correct means the analysis method successfully identifies all the precursor phases without any spurious phases. Blue bars indicate patterns scanned using the 2-minute (low-quality) program, while pink bars represent the 8-minute (medium-quality) scans. The top of each bar shows the number of correct predictions compared to the total number of patterns for that scan type. (c) Relationship between the R_{wp} values from Rietveld refinement using groundtruth phases and Dara's top result (represents the solution of lowest R_{wp} that Dara can find). The ground-truth phases are the precursor phases added during sample preparation. Blue and pink dots correspond to 2-minute and 8-minute scans, respectively. (d) Dara's runtime per pattern as a function of the number of reference phases in the database. Blue and pink dots represent 2-minute and 8-minute scans, respectively. Dashed horizontal lines mark 2 and 8 minutes on the time axis, corresponding to the measurement time to obtain these patterns in the diffractometer. The time is measured on a workstation with Intel(R) Core(TM) i9-10920X CPU @ 3.50GHz.

We compare Dara's phase identification performance against Jade,³⁹ a widely used commercial software for powder XRD analysis that integrates both phase identification and basic

full-profile fitting. Parameters for both methods are described in Supplementary Note ??.

We count the number of correct patterns identified by each analysis method, where a pattern is considered “correct” if the result includes all precursor phases added when preparing the sample and excludes any spurious phases. One exception is made for the pattern composed of 30 wt% NiO, 30 wt% Bi₂O₃, and 40 wt% Li₂CO₃. In this case, both Jade and Dara detect a minor Bi₂O₂(CO₃) phase, present at approximately 3 wt%, which may result from a reaction between Bi₂O₃ and Li₂CO₃ or CO₂ in the air during mixing. Although this phase exhibits distinct peaks in the XRD pattern, we choose to exclude it from our evaluation due to the lack of supporting evidence beyond the diffraction data. Results are shown in Figure 2(b).

For the 2-minute scans, Jade misclassifies 4 out of 20 patterns: either missing a phase or introducing spurious ones. Dara misclassifies only two patterns, both of which result from missing a phase. For the higher-quality 8-minute scans, both methods show improved accuracy due to better signal-to-noise ratios. In this case, Dara still outperforms Jade, identifying all 20 patterns successfully, while Jade still fails in 2 cases. Detailed analysis of Dara’s and Jade’s phase identification results for all samples is provided in Supplementary Note ??.

To further understand the fitness of the phases selected by Dara, we plot the weighted profile residuals (R_{wp}) of Dara’s top result (the result with the lowest R_{wp}) against the R_{wp} produced from human-performed Rietveld refinement. For each precursor used in constructing the dataset, the groundtruth reference phase is carefully hand-picked by humans from the structural database and a very good fit is obtained with the XRD pattern of the single precursor, as shown in Supplementary Note ??. Hence, these reference phases are considered sufficiently similar that they can be used as an approximation to the actual phase in the sample (groundtruth). The ICSD IDs of these phases are listed in Supplementary Table ??. Figure 2(c) compares the R_{wp} of refining with the known groundtruth phases versus Dara’s top-ranked refinement outcome, as listed in Supplementary Table ??. Across the dataset, Dara consistently returns solutions with $R_{wp} < 10\%$, reflecting good-quality fits and generally aligns with the groundtruth results. However, in some cases, the groundtruth R_{wp} is

lower than Dara’s. This discrepancy can arise from the two-stage nature of our approach and the varied refinement parameters used at different stages, as described in the *Method* section. After phase searching, Dara performs the final refinement using only the representative phases with the highest Figure of Merit (FoM). In some cases, the phases that fit best during the constraint search stage are not the ones that would fit best for the final refinement, where Rietveld refinement is allowed to adjust more parameters to achieve a better fit. Even so, the selected phases can still be regarded as valid matches to the experimental pattern, as the differences primarily arise from variations in peak intensity rather than the appearance or disappearance of specific peaks. These cases underscore the ambiguity in XRD analysis: multiple, subtly different phases can produce comparably good fits under different refinement parameters, and distinguishing between them often requires careful inspection and, in some cases, additional characterization.

Finally, we evaluate Dara’s use of computational resources. Different from Jade’s fast, heuristic search-match algorithm, which completes in seconds, Dara relies on a more exhaustive and detailed evaluation of phases that involve hundreds of Rietveld refinements for a single pattern. This approach admittedly requires more computational time but is better suited to handling complex, multi-phase patterns. Figure 2(d) shows that runtime scales with the number of reference phases. Dara processes 2-minute scans (blue dots) slightly faster than 8-minute scans (pink dots), mainly due to the fewer 2-theta angular steps measured in a shorter scan, which reduces the refinement time. This is because the refinement requires the computation of intensity at every angular point in the pattern to calculate the error. Several runtime optimization strategies are implemented in Dara: (1) an integrated heuristic search-match step to filter unlikely phases before committing to a full-profile Rietveld refinement that can be time-consuming; (2) grouping of structurally similar phases to avoid redundant refinements on XRD-equivalent variants (e.g., doped or vacancy-modified forms); and (3) parallel execution using the Ray framework,³² enabling scaled-up deployment on multi-node computing clusters. As a result, the typical runtime per pattern remains

shorter than the actual measurement time (2 or 8 minutes), marked as dashed horizontal lines in Figure 2(d).

Benchmarking on pairwise reaction product dataset

The second benchmark is designed to evaluate Dara’s performance on a dataset composed of the products of inorganic solid-state reactions. In these reactions, both precursors and products are typically solid powders. Due to the slow kinetics of solid-state diffusion, products often consist of off-stoichiometric solid solutions, unreacted precursors, and metastable intermediate phases.⁴⁰ These complex multiphase mixtures, frequently exhibiting significant variations in composition and structure, pose a major challenge for XRD analysis, despite XRD being a key technique for characterizing the outcomes of solid-state reactions. To this end, we construct a dataset comprising 20 samples from reactions between pairs of precursors chosen from 21 commonly used precursors, including oxides, carbonates, phosphates, and oxalates. Reaction temperatures are selected based on Tammann’s rule,⁴¹ which estimates the onset temperature of solid-state reactions to be roughly two-thirds of the lowest melting point among the precursors. Since this is often a somewhat low temperature, many reactions are incomplete with partially unreacted precursors. Such a multi-phase mixture tests Dara’s ability to make correct phase assignments for realistic samples that are often encountered in the exploratory phases of synthesis. The procedure to obtain the dataset is summarized in Figure 3(a). All experiments are performed in A-Lab, a fully automated synthesis platform equipped with robotic arms.¹⁰ Additional details on sample synthesis and XRD characterization are provided in Supplementary Note ???. The resulting XRD patterns are analyzed by a human expert using a typical XRD analysis and refinement approach, incorporating the PDF-5+ database and suggestions from Dara, to arrive at chemically reasonable interpretations (see Supplementary Note ?? for details). For benchmarking the automated tools, both Jade and Dara are also run in a fully unsupervised mode, without human intervention.




Figure 3 Preparation of the solid-state reaction dataset and benchmarking results. (a) Schematic illustration of the workflow for preparing the reaction dataset and approximating ground truth solutions through human expert evaluation. (b) Number of XRD patterns that have all peaks indexed in the analyses conducted by human experts, Dara, and Jade. (c) Comparison of the absolute lattice volume shifts after Rietveld refinement between the phases with the highest FoM scores (pink bars) and other phases with lower scores (blue bars) within each phase group found by Dara. A phase group refers to a set of phases that yield similar XRD patterns and are therefore expected to provide a comparable fit to the experimental peaks. The x-axis represents the absolute value of lattice volume shift (in %), and the y-axis shows the normalized frequency (with the sum of frequency set to 1). Pink bars on top indicate the volume shifts of the most appropriate phases selected by Dara in each phase group, based on a figure of merit (FoM) that incorporates both the quality factor ($1 - \rho$) and lattice shift (ΔU). Blue bars represent all other phases with lower FoM, which Dara did not choose to continue the search but considers as alternative phases to the FoM-selected phases in the result. The median for each histogram is shown as a dashed line in the plot. (d) Correlation between the R_{wp} values from human-expert Rietveld refinement and Dara's best-fit solution (i.e., the one yielding the lowest R_{wp} identified by Dara).

We first examine the accuracy of the phases identified by Dara. In solid-state reactions, it is common to form phases with compositions and lattice parameters that differ from those of the reference phases in the structure database. Sometimes, phases with a new structure that has not been included in the structure database can be encountered. In such cases, a good XRD pattern analyzer should still function effectively and identify all plausible known phases that can provide a good fit to the pattern. To evaluate Dara’s ability to handle such off-standard phases, we compare the phases identified by Dara, Jade, and the human expert across 20 XRD patterns with results shown in Figure 3(b). For each method, we count the number of patterns that have all peaks indexed. Only the top solution returned by Dara is considered. The human expert produces the most fully indexed patterns (16 out of 20), followed closely by Dara (15 out of 20), with four cases failing in the same manner as in the human’s analysis, potentially due to some unknown phases absent from the structure databases. The additional failure of Dara ($2\text{Fe}_3\text{O}_4 + 3\text{Y}_2\text{O}_3$ @ 1000 °C) occurs because it identifies the wrong polymorph of Y_2O_3 ($Fm\bar{3}m$ instead of the ground-state $Ia\bar{3}$), leading to a clearly unmatched peaks at around 20°. This is because the major phase, YFeO_3 , also has a minor peak at 20°, causing Dara’s peak matching algorithm to mistakenly treat the 20° peak as “matched” and deprioritize searching for phases that have a 20° peak. Jade, on the other hand, only manages to index 7 of the 20 patterns fully. We attribute this mainly to the fact that many phases formed in solid-state reactions differ from the reference phases in the structure database, which may make them difficult to capture via the search-match method. However, these phases can be captured by the full-profile fitting process, which models various XRD-related effects (e.g., sample displacement and broadened peaks) during the fitting phases. As for the precursor mixture dataset, we analyze the runtime of Dara for each pattern relative to the number of reference phases (Supplementary Figure ??). Since most samples in this benchmark contain fewer elements, there are fewer reference phases, and runtimes are generally lower than in the precursor mixture dataset. Most analysis workflows complete in under two minutes. Four samples required 2-8 minutes, while two

samples require more than 8 minutes, out of the total twenty patterns.

Dara can identify multiple phases that yield comparably good fits to a diffraction pattern, helping users recognize the range of possible interpretations before drawing conclusions from XRD data. For example, in the reaction $\text{Cr}_2\text{O}_3 + 2 \text{MnO}$ at 1100 °C, two spinel phases with distinct compositions are detected: MnCr_2O_4 and $\text{CrMn}_{1.5}\text{O}_4$. Due to the similar structure factor of Mn and Cr and the close lattice parameters ($a = 8.435\text{\AA}$ and $8.479\text{\AA}$, respectively), the XRD patterns of these two phases are hard to distinguish through a several-minute XRD scan on a lab X-ray diffractometer. For this example, Dara’s refinement indicates a spinel phase with a lattice parameter of $a = 8.454\text{\AA}$, suggesting a composition intermediate between the two spinels, potentially a Mn:Cr ratio of 1:1, which is indeed the Mn:Cr ratio in the precursor of this sample. This ability to identify potential fits is particularly useful for analyzing the reaction products when attempting to synthesize materials predicted by computational screening. By including the computed structure in the reference phase set, Dara is able not only to tell if the computed structure has a good match to the pattern but also to list all the known structures that have a good match in the structure database. If both the computed structure and the known phases can fit equally well to the pattern, Dara will catch the user’s attention and suggest additional analysis (e.g., SEM/EDS) to verify elemental compositions.

In Dara’s search, phases are grouped based on the similarity of their diffraction patterns. Phases within the same group are treated as effectively indistinguishable by XRD and can produce a similar fit to the given experimental pattern. From each group, a representative phase is selected using a figure of merit (FoM), adapted from Lutterotti et al.²⁴ This FoM, defined in Equation 1 (see Methods), combines a fitness term ($1 - \rho$), which evaluates the overall fit quality, and a lattice shift term (ΔU), which measures the change in lattice parameters a , b , and c before and after the refinement. To demonstrate the effectiveness of this selection criterion, we computed the lattice volume shifts during refinement for both top-FoM and lower-FoM phases in each phase group (Figure 3(c)). The top-FoM phases

(pink bars) exhibit smaller shifts, with a median of 0.15%, meaning the lattice volume is adjusted by only 0.15% during Rietveld refinement to align with the experimental peak positions. In contrast, the lower-FoM phases (blue bars) show larger shifts, with a median of 0.35%. This difference highlights that the FoM score reflects not only the fit quality (measured by quality factor $1 - \rho$) but also the changes in the lattice parameters, which often indicates compositional variation between the actual and reference phases, which is a common product in the reaction products of the solid-state reactions, primarily due to the formation of off-stoichiometric phases. This FoM prioritizes reference phases with lattice parameters more closely matching those of the sample, which can aid in estimating the composition of solid solutions present. A similar analysis of the precursor mixture dataset (Supplementary Figure ??) reveals the same trend. The top-FoM phases have a median lattice volume shift of 0.14%, while the lower-FoM phases show a median of 0.40%. Compared with the pairwise reaction dataset, the top-FoM phases in the precursor mixture dataset display slightly smaller shifts, whereas the lower-FoM phases deviate more. This is because the precursors used in the precursor mixture dataset have often been thoroughly characterized in the structure database and reported with larger ranges in lattice parameters, primarily due to the different synthesis methods and measurement conditions. Typically, at least one reference phase closely matches the lattice parameters of the sample, resulting in a slight lattice shift for the top-FoM phase. Other reported phases may require larger lattice adjustments to align with the experimental peaks, yet still achieve a comparable fit owing to their nearly identical crystal structures. As a result, Dara regards these phases as equally valid matches to the diffraction pattern. Distinguishing them requires more dedicated analysis of the XRD pattern or additional characterization.

Finally, we compare refinement outputs from Dara and human experts. Figure 3(d) shows a scatter plot of final R_{wp} values: Dara’s top (x-axis) versus the expert’s (y-axis). The points generally follow the $y = x$ line, indicating comparable fitting quality. Interestingly, human-derived R_{wp} values are often slightly higher, particularly in higher- R_{wp} cases. This may result

from the different refinement strategies: human experts emphasize physical interpretability and carefully adjust refinement parameters step by step; Dara employs a consistent, general-purpose refinement setup to achieve a good fit while minimizing the risk of overfitting. Although Dara’s Rietveld refinements are not intended for detailed microstructural analysis (e.g., grain size, strain, or site occupancy), they can provide reliable phase identification (indicated by the peak fitness) and estimation of phase fractions. Additionally, they can serve as excellent starting points for downstream, more detailed analysis to obtain information such as grain size, strain, and atomic occupancy from the XRD patterns.

Characterizing the ambiguity of XRD patterns

In practice, it is often hard to map the XRD pattern deterministically to the crystal structures due to the presence of multiple phases, instrument resolution limitations, peak overlap, and measurement noise.⁴² Multiple phases, with different compositions or slight variations in atomic arrangement, can produce almost indistinguishable XRD patterns. The limited resolution of laboratory X-ray diffractometers further amplifies this challenge. For example, diffraction peaks often overlap with one another or become obscured by background noise, making it challenging to identify individual phases confidently. In such cases, XRD alone will not be adequate to disambiguate these possibilities. Despite these limitations, lab-based XRD remains one of the most widely used tools for characterizing inorganic crystal samples.

We illustrate this challenge using a sample synthesized in our lab, shown in Figure 4. More details about sample preparation are described in Supplementary Note ???. It is the product of a solid-state reaction involving six elements: Li, Na, Al, Si, Co, and O. Due to the use of a Cu anode X-ray source and the presence of Co in the sample, significant background arises from secondary fluorescence, increasing the noise level. This results in a lower signal-to-noise ratio, making the XRD pattern more ambiguous for analysis. Nonetheless, although non-ideal for analysis, such patterns are typical of those encountered when analyzing real-world samples.

As shown in Figure 4, Dara identifies four groups of phase combinations (solutions) for this sample, each consisting of three phases. The R_{wp} for all four solutions ranges closely between 2.20% and 2.33%. Two major phases are consistently present across all four solutions. The first is a nepheline-type phase with a composition close to NaAlSiO_4 , for which Dara finds 11 matching reference phases in the structure database. The second is a series of solid solutions spanning the LiCoO_2 – LiAlO_2 tie line. These phases share similar lattice parameters (LiAlO_2 , $R\bar{3}m$: $a = 2.800\text{\AA}$, $c = 14.216\text{\AA}$; LiCoO_2 , $R\bar{3}m$: $a = 2.816\text{\AA}$, $c = 14.054\text{\AA}$) and can form across a wide range of compositions.^{43,44} Given this, it is difficult to determine the exact composition solely from the XRD measurement. While the two major phases remain the same across the four solutions, the third, minor phase picked by Dara varies and contains structurally and compositionally distinct phases: SiO_2 ($P3_221$), $\text{Co}_{11}\text{O}_{16}/\text{Co}_2\text{SiO}_4$ ($Fd\bar{3}m$) (7 matched phases), Al_2CoO_4 ($Fd\bar{3}m$) (25 matched phases), and NaCo_3O_4 . Since they have dissimilar diffraction patterns, these phases are grouped into four solution groups. For example, despite compositional differences between $\text{Co}_{11}\text{O}_{16}$ and Co_2SiO_4 , they are grouped due to their similar XRD patterns. As illustrated in the bottom row of Figure 4, each phase has its distinct diffraction peaks, but all exhibit a prominent peak that can be matched to a peak in the pattern at around 36.5° , as flagged with an orange triangle in Figure 4. However, due to substantial peak overlap between these minor phases and the major phases, it is challenging to determine which one is the actual phase that contributes to that peak. Detailed phase lists for each solution and composition groupings are provided in Supplementary Note ??.

Discussion

Dara has been deployed in both the autonomous laboratory (A-Lab)¹⁰ and standard research lab environments at Lawrence Berkeley National Laboratory. Through an internal web interface and API, Dara offers a user-friendly interface for automatically analyzing the




Figure 4 Example of multiple phase solutions identified by Dara for an experimental solid-state reaction sample. The raw XRD pattern and its corresponding chemical system are supplied to Dara. After searching, four solutions are found to fit the pattern similarly well, all of which contain three phases. The calculated patterns for each phase are displayed in the corresponding boxes in the plot. Phases 1 and 2 are shared across all the solutions, which are groups of NaAlSiO₄ with Nepheline structure (11 phases) and LiCoO₂-LiAlO₂ solid solutions (35 phases), respectively. Phase 3, however, includes four possible phases that differ greatly in structure/composition: SiO₂ (ICSD #155249), Co₁₁O₁₆/Co₂SiO₄ structure family (7 phases), Al₂CoO₄ structure family (25 phases), and NaCo₃O₄ (ICSD #163993), indicating that further compositional characterization may be necessary. The common peak at around 36.5° is marked with an orange triangle in the plot.

XRD data. As of July 13, 2025, it has processed 2,453 unique XRD patterns from our research group and external collaborators. Screenshots of the internal application are shown in Figure 5, including submission, overview, result details, and refinement plot visualization. This browser-based application makes Dara accessible to users who are not familiar with programming and XRD analysis.

Statistical analysis was performed on all Dara searches conducted through this web-based platform, as summarized in Figure 6. The first metric we examine is runtime. We find that Dara’s runtime scales with the number of reference phases considered during the search (Figure 6(a)). Since the uploaded patterns span a variety of chemical systems and the number of elements, runtimes vary accordingly. The median number of reference phases is 281 phases for one pattern, and the median runtime per pattern is 88.9 seconds, faster than a typical XRD data collection in a laboratory setting. Longer runtimes typically occur when the search space includes many reference phases, which can result from a chemical system with numerous elements or one that contains a large number of reference phases. In such cases, Dara may take several hours to complete an analysis. However, this can be significantly accelerated by deploying Dara on high-performance computing (HPC) clusters. With its parallel tree search implementation, Dara efficiently utilizes multiple cores/nodes to shorten processing times. In terms of solution quality, Dara successfully identifies at least one solution for most patterns. Figure 6(b) shows the distribution of R_{wp} . Dara achieves a median R_{wp} of 5.85%, with 78.1% of patterns yielding values below 10%. Although R_{wp} alone does not fully determine fit quality,⁴⁵ these statistics suggest that Dara performs robustly on real experimental lab data and provides a reliable starting point for human-guided phase identification and refinement.

Over several months of running this web application, we have observed some natural limitations that are also common to other XRD analysis software. The first is that the performance of Dara is highly related to the coverage of the experimental structure database in the chemical system of interest. For example, many structures have been well-characterized




Figure 5 Screenshots of the Dara web interface. (a) Analysis job submission page, where users input the pattern, elements that can exist in the pattern, and diffractometer information. (b) Overview page for viewing the status of and accessing each analysis job. (c) Result detail page with a summary of the job's outcome, including analysis parameters, the best result's R_{wp} , and the most probable phases. (d) Result detail page with all solutions and phases found by Dara. (e) Interactive plot to visualize the refinement produced by Dara.




Figure 6 Statistical analysis of patterns running on our internal Dara web-based application at Lawrence Berkeley National Laboratory. (a) The relationship between the runtime of analyzing a pattern and the number of reference phases. The vertical and horizontal dashed pink lines represent the median value for runtime and number of reference phases, respectively. The runtime is measured on a workstation with Intel(R) Core(TM) i9-10920X CPU @ 3.50GHz. (b) Histogram of the best R_{wp} values of each pattern obtained at the end of Dara search. The pink vertical line represents the median R_{wp} .

and deposited into structure databases within the $\text{Na}_2\text{O-Al}_2\text{O}_3\text{-SiO}_2$ chemical space. If one tries to analyze a pattern that contains Na, Al, Si, and O, it is more likely that Dara will find many solutions with various alternative phases. This is because the structures in this chemical space have been thoroughly studied, and numerous structures with minor structural modifications exist. On the other hand, if the sample contains phases that are not present in the structure database, Dara is likely to return either no result or phases that poorly fit the pattern. For instance, disordered rocksalt structures composed of Li-Mn-Ti-O have been extensively studied as promising cathode materials for next-generation lithium-ion batteries.⁴⁶ However, since no reference structure for these materials exists in the ICSD database, Dara fails to identify the Li-Mn-Ti-O phase when analyzing patterns from this composition space.

For powder XRD analysis, especially when using a lab diffractometer, it is often impossible to fully resolve the structures. This can be attributed to the fact that lab-diffractometer-based powder XRD cannot provide sufficient information about composition and atomic positions, especially when the sample contains a mixture of multiple phases. To solve a pattern, it must be matched to known structures characterized by other methods. In re-

cent years, the advancement of machine learning, particularly the emergence of generative models for crystal structures,^{47,48} has provided a viable approach to solving this problem. By conditioning the crystal generation on the fitness of XRD patterns and optionally the residual force in the lattice, one can obtain structures that can fit the XRD pattern while being physically meaningful.^{49–51} However, these models can only handle samples dominated by a single ordered phase, while real samples are often mixtures of known and unknown phases, potentially with disorder. One possible solution is to use phase analysis software, such as Dara, first to identify the known phases and then pass the unidentified peaks to crystal generation tools to search for a feasible structure resolution.

Another challenge is in assessing the chemical feasibility of specific phases. As Dara’s tree search is designed to identify as many plausible phases as possible from the structural database, it often returns multiple phase combinations that fit the XRD pattern, some of which are unlikely to be present given the sample’s synthesis conditions. For example, Dara may identify elemental metals such as lithium or sodium as having a good fit, even though these metals are highly reactive under ambient conditions where XRD measurements are typically performed. This occurs because elemental metals often have high-symmetry crystal structures, resulting in simple and strong diffraction peaks that can easily match those of nearby peaks in the measured pattern. While human experts often rely on this information to decide which phases to test for. However, automated XRD analysis tools typically do not incorporate knowledge of synthesis chemistry or stability, making it challenging to filter out chemically implausible phases without excluding valid ones. To address this limitation, information about the sample’s preparation conditions is essential. One possible improvement is to incorporate thermodynamic details into the analysis. For example, tools like reaction network analysis⁵² can be used to predict the likelihood of a phase forming under given reaction conditions. By considering the relative energies of reference phases, this approach can effectively eliminate thermodynamically unlikely phases from the analysis. Additionally, in recent years, large language models (LLMs) have emerged as promising tools for solving sci-

entific problems in ways that mimic human reasoning.^{53,54} It is therefore possible to develop an LLM-based filter that leverages synthesis and chemistry knowledge to eliminate phases that are highly unlikely to exist in a given sample. This approach would enhance the chemical interpretability of XRD solutions by integrating domain knowledge into the automated analysis, as well as the information obtained from other characterization methods.

Conclusion

In this work, we present the design and performance evaluation of Dara, a tool for the automated analysis of powder XRD patterns. Leveraging a tree search algorithm for phase identification, a peak-matching algorithm for rapid identification of promising phases, robust full-profile fitting with the BGMN refinement engine, and an intelligent grouping algorithm for identified phases, Dara is capable of handling realistic powder diffraction samples with various sample effects and multi-phase mixtures. It is designed to address the ambiguity issue in XRD-based characterization and to provide reliable phase analysis by explicitly generating and testing alternative hypotheses with a refinement program, similar to how a human expert might analyze a powder sample of unknown phases. By comparing the performance of Dara with other analysis software as well as a human expert, we demonstrate that (i) Dara can match the performance of a human in analyzing the phase components of an XRD pattern, and (ii) it can efficiently analyze the pattern within a reasonable time with the full-profile fitting. As more tools like Dara are developed and integrated into autonomous synthesis workflows, we envision a future where high-throughput, expert-level structural analysis becomes routine, accelerating self-driving materials discovery and characterization at scale.

Author Contributions

Y.F. and M.J.M: conceptualization, software, writing - original draft, writing - review and editing. C.L.R. and S.W.: Validation and investigation. G.C.: resources, supervision, methodology, project administration, writing - review and editing.

Conflicts of interest

There are no conflicts to declare.

Data availability

The source code for Dara is available at <https://github.com/idocx/dara>. The version of Dara used in this study is v1.0.0. The benchmark summaries (Dara, Jade, and human) for the precursor mixture and pairwise reaction product dataset are available as two spreadsheets. Human analysis of the pairwise reaction product dataset is available in another spreadsheet. The raw XRD patterns used for this study are available at <https://zenodo.org/records/17410051>.

Acknowledgement

This work was primarily funded by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, Materials Sciences and Engineering Division under contract no. DE-AC02-05-CH11231 (D2S2 programme, KCD2S2). C.L.R was funded by the U.S. Department of Energy, Office of Science, Basic Energy Sciences, Division of Materials Science, through the Office of Science Funding Opportunity Announcement (FOA) Number DE-FOA-0002676: Chemical and Materials Sciences to Advance Clean-Energy Technologies and Transform Manufacturing. S.W. was funded by the U.S. Department of Energy, Office of Science, Office

of Basic Energy Sciences, Materials Sciences and Engineering Division under contract No. DE-AC02-05-CH11231 (MINES project: The Science of Direct MINeral to Energy Storage Synthesis, FWP: FP00014914). S.W. was supported in part by the Jane Lewis Fellowship at UC Berkeley.

The authors thank Anubhav Jain (LBL), Lauren Walters (UC Berkeley), Bernardus Rendy (UC Berkeley), Tim Kodalle (LBNL), Guilhem Dezanneau (CNRS), Nobumichi Tamura (LBNL), Adam Corrao (BNL), Amalie Trewartha (TRI), Yan Zeng (FSU), and Olympia Dartsis (LBNL) for their testing and feedback. The authors would also like to express their special thanks to Nicola Döbelin (RMS Foundation) and Reinhard Kleeberg (TU Bergakademie Freiberg) for guidance and permission on using BGMN refinement software as Dara’s backend. Finally, we gratefully acknowledge our colleagues and peers for discussion and constructive feedback on the design of robust automated XRD analysis tools, which has greatly strengthened the development of this work.

References

- (1) Waseda, Y.; Matsubara, E.; Shinoda, K. *X-ray diffraction crystallography: introduction, examples and solved problems*; Springer Science & Business Media, 2011.
- (2) Ward, L.; Michel, K.; Wolverton, C. Automated crystal structure solution from powder diffraction data: Validation of the first-principles-assisted structure solution method. *Physical Review Materials* **2017**, *1*, 063802.
- (3) Narayan, A.; Bhutani, A.; Rubeck, S.; Eckstein, J. N.; Shoemaker, D. P.; Wagner, L. K. Computational and experimental investigation for new transition metal selenides and sulfides: the importance of experimental verification for stability. *Physical Review B* **2016**, *94*, 045105.
- (4) Shoemaker, D. P.; Hu, Y.-J.; Chung, D. Y.; Halder, G. J.; Chupas, P. J.; Soderholm, L.; Mitchell, J.; Kanatzidis, M. G. In situ studies of a platform for metastable inorganic crystal growth and materials discovery. *Proceedings of the National Academy of Sciences* **2014**, *111*, 10922–10927.

- (5) McCusker, L.; Von Dreele, R.; Cox, D.; Louër, D.; Scardi, P. Rietveld refinement guidelines. *Journal of Applied Crystallography* **1999**, *32*, 36–50.
- (6) Jain, A.; Ong, S. P.; Hautier, G.; Chen, W.; Richards, W. D.; Dacek, S.; Cholia, S.; Gunter, D.; Skinner, D.; Ceder, G.; others Commentary: The Materials Project: A materials genome approach to accelerating materials innovation. *APL materials* **2013**, *1*.
- (7) Zagorac, D.; Müller, H.; Ruehl, S.; Zagorac, J.; Rehme, S. Recent developments in the Inorganic Crystal Structure Database: theoretical crystal structure data and related features. *Journal of applied crystallography* **2019**, *52*, 918–925.
- (8) Kabekkodu, S. N.; Dosen, A.; Blanton, T. N. PDF-5+: a comprehensive powder diffraction file™ for materials characterization. *Powder Diffraction* **2024**, 1–13.
- (9) Gražulis, S.; Daškevič, A.; Merkys, A.; Chateigner, D.; Lutterotti, L.; Quirós, M.; Serebryanaya, N. R.; Moeck, P.; Downs, R. T.; Le Bail, A. Crystallography Open Database (COD): an open-access collection of crystal structures and platform for world-wide collaboration. *Nucleic Acids Research* **2012**, *40*, D420–D427.
- (10) Szymanski, N. J.; Rendy, B.; Fei, Y.; Kumar, R. E.; He, T.; Milsted, D.; McDermott, M. J.; Gallant, M.; Cubuk, E. D.; Merchant, A.; others An autonomous laboratory for the accelerated synthesis of novel materials. *Nature* **2023**, *624*, 86–91.
- (11) Lunt, A. M.; Fakhruddin, H.; Pizzuto, G.; Longley, L.; White, A.; Rankin, N.; Clowes, R.; Alston, B.; Gigli, L.; Day, G. M.; others Modular, multi-robot integration of laboratories: an autonomous workflow for solid-state chemistry. *Chemical Science* **2024**, *15*, 2456–2463.
- (12) Chen, J.; Cross, S. R.; Miara, L. J.; Cho, J.-J.; Wang, Y.; Sun, W. Navigating phase diagram complexity to guide robotic inorganic materials synthesis. *Nature Synthesis* **2024**, 1–9.
- (13) Yotsumoto, Y.; Nakajima, Y.; Takamoto, R.; Takeichi, Y.; Ono, K. Autonomous robotic experimentation system for powder X-ray diffraction. *Digital Discovery* **2024**,
- (14) Hanawalt, J.; Rinn, H.; Frevel, L. Chemical analysis by X-ray diffraction. *Industrial & Engineering Chemistry Analytical Edition* **1938**, *10*, 457–512.
- (15) Smith, D.; Gortner, S. Powder diffraction program information 1990 program list. *Journal of applied crystallography* **1991**, *24*, 369–402.

- (16) Lin, S.-F.; Lu, X.-J.; Zheng, W.-F.; Zhang, J.-B. A novel search/match system for X-ray powder diffraction data. *Chemometrics and intelligent laboratory systems* **1993**, *20*, 85–91.
- (17) Dinnebier, R. E.; Billinge, S. J. *Powder diffraction: theory and practice*; Royal society of chemistry, 2015.
- (18) Oviedo, F.; Ren, Z.; Sun, S.; Settens, C.; Liu, Z.; Hartono, N. T. P.; Ramasamy, S.; DeCost, B. L.; Tian, S. I.; Romano, G.; others Fast and interpretable classification of small X-ray diffraction datasets using data augmentation and deep neural networks. *npj Computational Materials* **2019**, *5*, 60.
- (19) Lee, J.-W.; Park, W. B.; Lee, J. H.; Singh, S. P.; Sohn, K.-S. A deep-learning technique for phase identification in multiphase inorganic compounds using synthetic XRD powder patterns. *Nature communications* **2020**, *11*, 86.
- (20) Maffettone, P. M.; Banko, L.; Cui, P.; Lysogorskiy, Y.; Little, M. A.; Olds, D.; Ludwig, A.; Cooper, A. I. Crystallography companion agent for high-throughput materials discovery. *Nature Computational Science* **2021**, *1*, 290–297.
- (21) Szymanski, N. J.; Bartel, C. J.; Zeng, Y.; Tu, Q.; Ceder, G. Probabilistic deep learning approach to automate the interpretation of multi-phase diffraction spectra. *Chemistry of Materials* **2021**, *33*, 4204–4215.
- (22) Zhang, S.; Cao, B.; Su, T.; Wu, Y.; Feng, Z.; Xiong, J.; Zhang, T.-Y. Crystallographic phase identifier of a convolutional self-attention neural network (CPICANN) on powder diffraction patterns. *IUCrJ* **2024**, *11*, 634.
- (23) Suzuki, Y.; Hino, H.; Hawai, T.; Saito, K.; Kotsugi, M.; Ono, K. Symmetry prediction and knowledge discovery from X-ray diffraction patterns using an interpretable machine learning approach. *Scientific reports* **2020**, *10*, 21790.
- (24) Lutterotti, L.; Pilliere, H.; Fontugne, C.; Boullay, P.; Chateigner, D. Full-profile search–match by the Rietveld method. *Journal of applied crystallography* **2019**, *52*, 587–598.
- (25) Chang, M.-C.; Ament, S.; Amsler, M.; Sutherland, D. R.; Zhou, L.; Gregoire, J. M.; Gomes, C. P.; van Dover, R. B.; Thompson, M. O. Probabilistic Phase Labeling and Lattice Refinement for Autonomous Material Research. *arXiv preprint arXiv:2308.07897* **2023**,

- (26) Leeman, J.; Liu, Y.; Stiles, J.; Lee, S. B.; Bhatt, P.; Schoop, L. M.; Palgrave, R. G. Challenges in High-Throughput Inorganic Materials Prediction and Autonomous Synthesis. *PRX Energy* **2024**, *3*, 011002.
- (27) Holder, C. F.; Schaak, R. E. Tutorial on powder X-ray diffraction for characterizing nanoscale materials. 2019.
- (28) Gražulis, S.; Chateigner, D.; Downs, R. T.; Yokochi, A. F. T.; Quirós, M.; Lutterotti, L.; Manakova, E.; Butkus, J.; Moeck, P.; Le Bail, A. Crystallography Open Database – an open-access collection of crystal structures. *Journal of Applied Crystallography* **2009**, *42*, 726–729.
- (29) Ong, S. P.; Richards, W. D.; Jain, A.; Hautier, G.; Kocher, M.; Cholia, S.; Gunter, D.; Chevrier, V. L.; Persson, K. A.; Ceder, G. Python Materials Genomics (pymatgen): A robust, open-source python library for materials analysis. *Computational Materials Science* **2013**, *68*, 314–319.
- (30) Doebelin, N.; Kleeberg, R. Profex: a graphical user interface for the Rietveld refinement program BGMN. *Applied Crystallography* **2015**, *48*, 1573–1580.
- (31) Bergmann, J.; Taut, T. Rietveld Analysis Program BGMN, 4th revised edition. Dr. J. Bergmann: Dresden, Germany, 2005; Copyright © 1996, 1998, 1999, 2005 by J. Bergmann; BGMN is a German registered trademark of Dr. Jörg Bergmann.
- (32) Moritz, P.; Nishihara, R.; Wang, S.; Tumanov, A.; Liaw, R.; Liang, E.; Paul, W.; Jordan, M. I.; Stoica, I. Ray: A distributed framework for emerging AI applications. CoRR abs/1712.05889 (2017). *arXiv preprint arXiv:1712.05889* **2017**,
- (33) Bergmann, J. EFLECH/INDEX—another try of whole pattern indexing. *Z. Kristallogr. Suppl* **2007**, *26*, 197–202.
- (34) Ziegel, E. R. The elements of statistical learning. 2003.
- (35) Pedregosa, F. et al. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* **2011**, *12*, 2825–2830.
- (36) Bergmann, J.; Friedel, P.; Kleeberg, R. BGMN—a new fundamental parameters based Rietveld program for laboratory X-ray sources, its use in quantitative analysis and structure investigations. *CPD Newsletter* **1998**, *20*, 5–8.

- (37) Taut, T.; Kleeberg, R.; Bergmann, J. Seifert Software: The new Seifert Rietveld program BGMN and its application to quantitative phase analysis. *Materials Structure* **1998**, *5*, 57–66.
- (38) Jenks, G. F. The data model concept in statistical mapping. *International yearbook of cartography* **1967**, *7*, 186–190.
- (39) Materials Data, Inc. JADE Pro. <https://materialsdata.com/>, 2019; Materials Data, Inc., Livermore, CA, USA.
- (40) McDermott, M. J.; McBride, B. C.; Regier, C. E.; Tran, G. T.; Chen, Y.; Corrao, A. A.; Gallant, M. C.; Kamm, G. E.; Bartel, C. J.; Chapman, K. W.; others Assessing thermodynamic selectivity of solid-state reactions for the predictive synthesis of inorganic materials. *ACS Central Science* **2023**, *9*, 1957–1975.
- (41) Merkle, R.; Maier, J. On the tammann–rule. *Zeitschrift für anorganische und allgemeine Chemie* **2005**, *631*, 1163–1166.
- (42) Schreiner, W. Systematic and random powder diffractometer errors relevant to phase identification. *Norelco Rep.* **1982**, *29*, 42.
- (43) Khan, M. N.; Bashir, J. Synthesis and structural refinement of $\text{LiAl}_x\text{Co}_{1-x}\text{O}_2$ system. *Materials research bulletin* **2006**, *41*, 1589–1595.
- (44) Buta, S.; Morgan, D.; Van der Ven, A.; Aydinol, M.; Ceder, G. Phase Separation Tendencies of Aluminum-Doped Transition-Metal Oxides ($\text{LiAl}_{1-x}\text{M}_x\text{O}_2$) in the $\alpha\text{-NaFeO}_2$ Crystal Structure. *Journal of The Electrochemical Society* **1999**, *146*, 4335.
- (45) Toby, B. H. R factors in Rietveld analysis: How good is good enough? *Powder diffraction* **2006**, *21*, 67–70.
- (46) Clément, R.; Lun, Z.; Ceder, G. Cation-disordered rocksalt transition metal oxides and oxyfluorides for high energy lithium-ion cathodes. *Energy & Environmental Science* **2020**, *13*, 345–373.
- (47) Zeni, C.; Pinsler, R.; Zügner, D.; Fowler, A.; Horton, M.; Fu, X.; Wang, Z.; Shysheya, A.; Crabbé, J.; Ueda, S.; others A generative model for inorganic materials design. *Nature* **2025**, *639*, 624–632.
- (48) Xie, T.; Fu, X.; Ganea, O.-E.; Barzilay, R.; Jaakkola, T. Crystal diffusion variational autoencoder for periodic material generation. *arXiv preprint arXiv:2110.06197* **2021**,

- (49) Guo, G.; Saidi, T. L.; Terban, M. W.; Valsecchi, M.; Billinge, S. J.; Lipson, H. Ab initio structure solutions from nanocrystalline powder diffraction data via diffusion models. *Nature Materials* **2025**, 1–9.
- (50) Riesel, E. A.; Mackey, T.; Nilforoshan, H.; Xu, M.; Badding, C. K.; Altman, A. B.; Leskovec, J.; Freedman, D. E. Crystal structure determination from powder diffraction patterns with generative machine learning. *Journal of the American Chemical Society* **2024**, *146*, 30340–30348.
- (51) Li, Q.; Jiao, R.; Wu, L.; Zhu, T.; Huang, W.; Jin, S.; Liu, Y.; Weng, H.; Chen, X. Powder diffraction crystal structure determination using generative models. *Nature Communications* **2025**, *16*, 7428.
- (52) McDermott, M. J.; Dwaraknath, S. S.; Persson, K. A. A graph-based network for predicting chemical reaction pathways in solid-state materials synthesis. *Nature communications* **2021**, *12*, 3097.
- (53) Lei, G.; Docherty, R.; Cooper, S. J. Materials science in the era of large language models: a perspective. *Digital Discovery* **2024**, *3*, 1257–1272.
- (54) Jablonka, K. M.; Ai, Q.; Al-Feghali, A.; Badhwar, S.; Bocarsly, J. D.; Bran, A. M.; Bringuier, S.; Brinson, L. C.; Choudhary, K.; Circi, D.; others 14 examples of how LLMs can transform materials science and chemistry: a reflection on a large language model hackathon. *Digital discovery* **2023**, *2*, 1233–1250.