

Supplementary Information

Adaptive X-ray diffraction with automated phase identification

Nathan J. Szymanski^{1,2}, Christopher J. Bartel^{1,2}, Yan Zeng², Mouhamad Diallo^{1,2},
Haegyeom Kim², and Gerbrand Ceder^{1,2,*}

Information included:

Supplementary Notes 1, 2

Supplementary Figures 1-11

Supplementary Table 1

¹Department of Materials Science & Engineering, UC Berkeley, Berkeley, CA 94720, USA

²Materials Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

*Corresponding author

Supplementary Note 1.


The total measurement time required by adaptive XRD is controlled by two parameters: the confidence cutoff and the resampling threshold. The confidence cutoff dictates when more sampling is needed to identify a given phase, while the resampling threshold tells the diffractometer which 2θ to scan with increased resolution. Here, a scan with “increased resolution” refers to measurements performed with a high sampling density and long scan time (per step). When predictions from XRD-AutoAnalyzer have a confidence less than the specified cutoff, additional scans are proposed. For resampling of 2θ , scans are focused on regions with large differences between the CAMs of suspected phases (Methods section). A low resampling threshold leads to long measurement times, but also more accurate predictions since a greater portion of pattern is resampled with higher resolution. Our tests reveal that the F_1 -score begins to level off after an average measurement time of 10-15 min per pattern (Fig. 2), which corresponds to a resampling threshold of 25%. Slight improvements in accuracy can be observed upon decreasing the threshold further; however, this comes at the cost of substantially longer measurement time. Accordingly, to maintain a balance between speed and accuracy, a resampling threshold of 25% was used throughout all experimental tests described in the main text.

To choose an optimal confidence cutoff, we analyzed the distribution of confidence across simulated patterns in the Li-La-Zr-O space. These confidence values were divided into two cases: when XRD-AutoAnalyzer made correct versus incorrect classifications. As shown in **Supplementary Fig. 2a**, there is a clear distinction between these two distributions. On average, correct classifications are made with a high confidence of 81.5%, whereas incorrect classifications are made with a much lower confidence of 32.4%. Furthermore, except for a few outliers (< 1% of all patterns), all incorrect classifications are made with a confidence below 50%. This motivated our choice of 50% as a confidence cutoff that was used throughout all tests reported in the main text. Indeed, the use of a higher confidence cutoff does not lead to any noticeable improvement in the F_1 -score on simulated patterns, as shown in **Supplementary Fig. 2b**, because nearly all phases with a confidence greater than 50% are already correctly identified. Resampling such patterns only leads to an increased scan time, with little to no effect on which phases are predicted.


Supplementary Note 2.

As described in the Methods section, CAMs give information regarding which input features contribute most to the classification made by a CNN. Here we find that CAMs tend to be maximal in regions of 2θ that contain the largest peaks in a given XRD pattern. To illustrate this, we simulated ideal patterns (no artifacts) for all phases in the Li-La-Zr-O chemical space, and calculated the CAM associated with each phase using the trained CNN. In **Supplementary Fig. 3**, we show a parity plot that compares the position (2θ) of the largest peak in each pattern with the maximum of the CAM calculated for the corresponding phase. A clear correlation is observed between these two properties for most cases. There are some minor exceptions where a phase has two comparably large peaks that occur far apart in the XRD pattern – for example, Li_2CO_3 has two peaks with intensities that differ by only 3%, but positions that differ by $\sim 10^\circ$. The CAM associated with this phase is large around each of these two major peaks, but the maximum coincides with the one that has slightly weaker intensity.


To select which 2θ should be resampled with adaptive XRD, one option is to prioritize areas where the CAMs of suspected phases are large. However, as described in the main text, we hypothesize that this method is not optimal as it will always lead to increased sampling around the largest peaks in a diffraction pattern, which are already easy to detect using fast scans with low resolution. To confirm this hypothesis, we carried out tests with adaptive XRD on simulated patterns in the Li-La-Zr-O chemical space using two different methods: (1) prioritizing areas with maximal CAM of the most probable phase predicted by XRD-AutoAnalyzer, and (2) prioritizing areas with maximal *differences* between CAMs of the two most probable phases (Methods section). The results of these tests are displayed in **Supplementary Fig. 5**. As suspected, method (2) leads to more rapid convergence to a high F_1 -score. Accordingly, CAM differences are prioritized for resampling throughout all tests discussed in the main text.




Supplementary Figure 1. F₁-scores achieved by XRD-AutoAnalyzer when applied to simulated patterns in the Li-La-Zr-O chemical space. Each line represents data calculated using a different starting range for adaptive XRD. Because predictions are aggregated across all ranges below each given $2\theta_{\max}$, the choice of starting range has a significant influence on the final (ensemble) results. We find that a starting range of $2\theta = [10^\circ, 60^\circ]$ gives the highest F₁-score across all angles tested and is therefore used throughout all calculations discussed in the main text.




Supplementary Figure 2. **a**, Distributions of confidence given by XRD-AutoAnalyzer when applied to simulated patterns from materials in the Li-La-Zr-O chemical space. These are divided into correct and incorrect classifications to show the distinction in confidence. Each violin plot shows the density of samples with a given confidence value, while the boxes extend from the lower to upper quartile of the data. Black dots represent the average confidence given each case. **b**, The F_1 -score achieved by adaptive XRD on simulated patterns from materials Li-La-Zr-O chemical space, including mixtures with as many as three phases. Average scan times are also shown according to the right y-axis. Each dot represents calculations carried out with a different confidence cutoff, denoted by the x-axis.




Supplementary Figure 3. Parity plot showing the positions where maxima occur in the XRD patterns and CAMs calculated for each phase in the Li-La-Zr-O chemical space. A clear correlation is observed between these two properties, confirming the CAMs tend to be maximal near the largest peaks in XRD patterns. There are a few exceptions where comparably large diffraction peaks occur far apart in 2θ (*e.g.*, Li_2CO_3). In such cases, the CAM tends to be multi-modal, showing maxima around each of the largest peaks.




Supplementary Figure 4. A schematic highlighting the differences between individual and ensemble analyses. The individual approach makes a single prediction across the specified range of 2θ . In contrast, the ensemble approach makes several predictions using models separately trained on $10-60^\circ$, $10-70^\circ$, ..., $10-2\theta_{\max}$ (in this case, up to 90°) and aggregates those predictions using a confidence-weighted sum.




Supplementary Figure 5. F₁-scores achieved on patterns from materials in the Li-La-Zr-O chemical space when using a modified version of adaptive XRD (denoted CAM Max), where regions of 2θ with maximal CAM are prioritized for resampling. For comparison, we show results from the finalized version of adaptive XRD (denoted CAM Diff.), which prioritizes regions of 2θ with large *differences* between CAMs of suspected phases. The latter version is shown to reach convergence more quickly than the former, and therefore CAM differences are always prioritized during the tests described in the main text.




Supplementary Figure 6. Signal-to-noise ratios of the largest peaks in XRD patterns of Li_2CO_3 measured with a Bruker D8 Advance diffractometer with varied scan times per step. A fixed step size of $\Delta 2\theta = 0.02^\circ$ was used. The curved blue line is generated from the equation, $s/n = C\sqrt{t}$, where $C = 6.12$ is calculated *via* least-squares fitting of the observed data (assuming t in ms).




Supplementary Figure 7. An example illustrating how adaptive sampling was carried out on simulated XRD pattern. To mimic a fast experimental scan followed by slower resampling, high-resolution data was interpolated into the low-resolution pattern as shown in the lower plot.




Supplementary Figure 8. F_1 -scores achieved by XRD-AutoAnalyzer when applied to simulated patterns in the (top) Li-La-Zr-O and (bottom) Li-Ti-P-O spaces. Conventional results (black squares) were obtained on patterns with incrementally improved resolution over $2\theta = [10^\circ, 60^\circ]$, whereas adaptive results (blue circles) were found by resampling a subset of $2\theta \subseteq [10^\circ, 60^\circ]$ with high resolution. In contrast to **Fig. 2** included in the main text, a longer scan time of 40 min was used for the upper bound to show that such measurements (highlighted red) yield only marginal improvements in the F_1 -score relative to 30 min scans.



Supplementary Figure 9. The minimum weight fractions (%) of minority phases that can be detected by XRD-AutoAnalyzer when applied to mixtures characterized by conventional versus adaptive scans. Results are displayed for materials in the **a**, Li-La-Zr-O and **b**, Li-Ti-P-O chemical spaces.



Supplementary Figure 10. Two examples from the minority phase detection tests, showing how adaptive XRD improves the accuracy of phase identification on experimental patterns. **a**, In a mixture of ZrO_2 (90% wt.) and LiOH (10% wt.), LiOH is successfully detected by resampling 2θ from 29° to 38° with increased resolution to reveal several important peaks that would otherwise blend in with the background noise. **b**, In a mixture of Li_3PO_4 (84% wt.) and Li_2TiO_3 (16% wt.), Li_2TiO_3 is successfully detected by expanding 2θ beyond 60° to capture an additional peak that distinguishes it from Li_3PO_4 .



Supplementary Figure 11. An example illustrating how adaptive XRD can detect phases missed by conventional scans when applied *in situ* during the synthesis of LLZO. At 100 °C, high-resolution measurements are focused on ranges of 2θ that contain diffraction peaks associated with Li_2CO_3 . Several important peaks (diamond markers) are difficult to detect using data from fast scanning but become clearer in the adaptively sampled pattern, facilitating the detection of Li_2CO_3 .

Supplementary Table 1. References phases used to train and validate XRD-AutoAnalyzer.

Li-La-Zr-O		Li-Ti-P-O	
Formula	Space group no.	Formula	Space group no.
ZrO	225	LiTiPO ₅	62
Li ₇ La ₃ Zr ₂ O ₁₂	142	Li ₃ P	194
ZrO ₂	19	P ₂ O ₅	62
Li ₄ Zr ₃ O ₈	141	LiPO ₄	146
Li ₇ La ₃ Zr ₂ O ₁₂	230	LiTi ₂ O ₄	58
La ₂ O ₃	206	TiP ₂ O ₉	14
La ₂ O ₃	164	Ti(PO ₃) ₃	9
Zr ₃ O	167	Li ₄ TiO ₄	63
LiOH-H ₂ O	12	TiPO ₄	31
La(OH) ₃	176	Ti ₄ (PO ₅) ₃	43
Li ₂ ZrO ₃	15	TiP	194
Li ₂ CO ₃	15	Li ₈ PO ₃	221
La ₂ Zr ₂ O ₇	227	P ₂ O ₅	161
Li ₂ O	225	Li ₃ PO ₄	31
ZrO ₂	61	Ti ₃ PO ₂	164
Zr ₂ O	224	TiP	193
ZrO ₂	62	Li ₂ Ti ₆ O ₁₃	12
ZrO ₂	137	LiTi ₂ O ₄	12
LiO ₂	58	Li ₂ TiO ₃	15
Zr ₃ O	182	Li ₂ Ti ₃ O ₇	11
Li ₂ O ₂	194	TiP ₂ O ₇	205
LaOOH	11	LiTi ₂ O ₄	11
LiLaO ₂	14	Li ₂ O	225
ZrO ₂	14	LiTi ₂ O ₄	227
Li ₆ Zr ₂ O ₇	15	LiPO ₃	13
LiOH	129	TiPO ₄	11
ZrO ₂	225	Li ₂ Ti ₂ (PO ₄) ₃	60
ZrO	189	LiTiO ₂	141
		LiTiO ₂	227
		LiPO ₃	14
		TiO ₂	136
		LiTiPO ₅	2
		LiTi ₂ O ₄	31
		TiPO ₄	14
		Li ₄ P ₂ O ₇	14
		Li ₂ O	166
		LiTi ₂ (PO ₄) ₃	167
		P ₂ O ₅	43
		Ti ₅ (PO ₅) ₄	19
		Li ₄ P ₂ O ₇	2
		TiP ₂	62
		Li ₃ PO ₄	62
		LiTiP ₂ O ₇	4
		Li ₄ P ₄ O ₁₇	15
		TiO ₂	141

